

The background of the cover features a futuristic, abstract design. It consists of several translucent, overlapping panels that resemble digital screens or windows. These panels are filled with various types of data, including what look like code snippets, charts, and lists of numbers. The colors are primarily shades of blue, purple, and white, creating a high-tech, information-rich atmosphere. The overall effect is one of a complex digital environment or a sophisticated analytical tool.

Driving revenue growth through innovation

A how-to guide to monetizing GenAI and AI agents

Abde Tambawala
Alix Nepveux
Alexander Ammer

simon-kucher.com

SIMON 
KUCHER
Unlocking better growth

Contents

Welcome to the Agentic AI era	3
How to drive profitable revenue growth through AI innovation	8
Outcome-based pricing & packaging	21
How to adjust operations for agentic AI and GenAI	33
Turning into action	37
Authors	40

Welcome to the Agentic AI era

A year ago, most conversations were about large language models, foundation models, or copilots. These were still framed as assistive tools: things you prompt, things that augment human effort, things that sit alongside you. Now the narrative is shifting: not just tools that help, but agents that act. Instead of being reactive to prompts, agents can be proactive, goal-oriented, and autonomous across workflows.

The projections are staggering, almost doubling the size of the market in just four or five years. Agents promise to swallow entire categories of repetitive, operational work, serving as the connective tissue that sits between “ideas” and “outcomes.”

This shift to agents isn’t happening in a vacuum. It’s the next step in a much longer journey of how businesses adopt and value software. Each prior era reshaped not only the technology but also the mental model of what software is for and how it is paid for.

In the on-premise era, software was essentially infrastructure. Customers bought big systems, paid heavy upfront CAPEX, and installed them in their own servers. They were back-office enablers, not frontline drivers of growth. Scalability was limited, innovation cycles were slow, and value was captured mainly by the vendors who could sell big license deals to enterprises.


Then came the SaaS era, roughly from 2010 through today. This was the breakthrough that turned software into a utility customers could scale flexibly. Instead of massive upfront investment, they subscribed. Instead of running their own servers, they tapped the cloud. Software shifted from back-office to the center of business operations: CRM, ERP, HR, finance, marketing, productivity. It became the backbone of how companies run. And along the way, pricing matured around the subscription model: licenses, seats, tiers, usage.

“

The software market saw strong growth when moving from Perpetual licences to ‘software-as-a-service’; imagine how much it could expand shifting to ‘Services-as-a-software’, capturing a portion of the business services market on top of its current TAM

– [Alix Nepveux, Partner](#)

Now we’re on the cusp of the AI Agent era. What’s different here is that software is no longer just a tool. Agents can autonomously execute workflows, make recommendations, and, if they are sophisticated, take actions on their own. The growth projections show how transformative this is expected to be: the market could almost double in just four or five years as agents spread into sales, service, operations, compliance, healthcare, finance, and beyond.



This new era is also redefining value. In the on-premise era, value was tied to infrastructure. In the SaaS era, value was tied to access and scalability. In the AI Agent era, value will be tied to outcomes.

You may expect this paper to tell a tale of technology. However, we reframe this development as a business model story. Because as agents really do start mirroring or augmenting human workflows, then the way companies buy and pay for software will have to change with it.

Why does the AI agent era feel so different?

There's a reason why the market projections are so aggressive. Up to now, automation was mostly about workflows (rules, triggers, moving information from one place to another, etc.) However, cognitive work, the kind of thinking humans have traditionally done, was left untouched.

That's what agents are beginning to crack. They can support decision-making in areas such as sales outreach and customer engagement, tackle complex problem-solving tasks, and conduct sophisticated research and analysis.

The second driver is personalization at scale. Humans can only handle so much tailoring before it becomes overwhelming. AI agents, on the other hand, can customize interactions for thousands of customers or employees simultaneously, in real time. Every email, every product recommendation, every HR policy reminder, all personalized. This changes the nature of customer experience and employee engagement, opening up new expectations and new budgets.

The third factor is the data explosion. Enterprises are drowning in data they can't fully use. Legacy analytics tools help, but they're slow, require expertise, and often leave

insights buried. AI agents can chew through that data, find patterns, and surface actions in seconds. That makes the value proposition hard to ignore, because companies know they're sitting on underutilized assets.

Finally, there are new revenue streams. AI-native applications bring fresh monetization paths: usage-based models tied to consumption, commission-style models tied to performance, and embedded intelligence layered into existing SaaS platforms as premium upsells. Just like SaaS birthed an entirely new ecosystem of products and companies, the agent era is poised to create its own universe of monetization opportunities.

Addressing the payment disconnect

Modern employment is structured around renting time and skills through wages.



Compensation has traditionally been tied either to **time** (hourly pay), **performance** (commissions), or **scarce expertise** (high hourly rates for specialists).

The idea of compensating workers by the hour emerged during the Industrial Revolution in the late 18th and early 19th centuries. Beforehand, many workers were paid by the piece (piecework in agriculture or textiles) or through fixed contracts. As factories spread, employers needed a way to control and measure labor more consistently. The clock became central. Time itself was commodified, and hourly wages became the norm for lower- and mid-skilled labor – a model that persists today.

From an employer's mindset, this created a simple equation: more hours worked = more pay owed. Productivity could be adjusted by scheduling more or fewer hours. It's a transactional model that matches effort directly with compensation.

That same question now confronts the world of AI. As AI systems begin to take on human-like tasks, emulating or augmenting them, the old SaaS levers of “features and seats” start to look out of step.

Why do we think of humans in terms of hours and performance, but refrain from applying the same logic to machines? Nobody hires a person by asking, “List me every possible feature of your brain.” Instead, they ask, “What role can you play? What problems can you solve? How reliable are you? How much responsibility can I trust you with?”


The same will be true of AI.

This isn't to disown the SaaS packaging and pricing concepts that defined the last decade. In many ways, they were a breakthrough: built around what the software enabled users to do and translating feature lists into tangible business benefits and outcomes.

What's changing with AI agents is not the focus on outcomes, but the nature of how those outcomes are delivered. The line between "software capability" and "work execution" is blurring. Agents begin to deliver the benefits autonomously. The value, therefore, no longer lies only in what the agent can do, but in how much trust you can place in it to do it on its own, much like evaluating a human employee. And just as with human workers, that trust directly shapes the price you are willing to pay.

This requires a mindset-shift across the industry, moving away from viewing AI as a bundle of features or lines of code to be licensed, and instead treating it as a contributor whose value is defined by its role, level of autonomy, expertise, and the outcomes it delivers. Until that shift happens, monetization models will remain misaligned with the true potential of agentic AI.

Most software companies have already launched AI features, and almost all companies plan to launch AI features in the near future



What's striking is not who's building, it's how few are truly monetizing. Most launches remain incremental, tech-led add-ons that showcase capability rather than solving a clear customer problem. Vendors push out AI features because they can, not always because they should. And even when real value is created, many companies are unsure how to capture it commercially.

In other words, AI ambition is everywhere, but effective monetization is still rare.

Most AI providers face challenges pricing solutions

Many AI providers are still feeling their way through pricing. The technology has leapt forward faster than the business models have matured, and the tension is visible across the industry. We frequently observe that this challenge starts early in the process. GenAI and agentic solutions are often built without a clear definition of their commercial role and targets - whether they are meant to drive acquisition and usage, monetize incremental value, or improve overall profitability. When strategic intent is unclear, pricing decisions become a challenge, increasing the likelihood of either underpricing relative to value or structurally unprofitable usage.

Think about GPT products, which have taken the world by storm. When first launched, pricing wasn't the result of carefully modeled economics, rather an instinctive decision. Open AI's CEO, Sam Altman even admitted "I personally chose the price and thought we would make some money." This approach mirrored what many early AI product builders did: releasing straightforward pricing models, not only because they expected to generate some revenue, but also to unlock large-scale adoption as quickly as possible.

After this initial "simple pricing" phase, the market moved into a period of rapid experimentation. OpenAI, GitHub, Clay, and others began iterating at high velocity, testing new monetization structures, and shifting away from pure per-user pricing. Their varied approaches, from usage-based models to hybrid tiers, highlight just how experimental and dynamic AI pricing has become. Providers are trying to find a number that balances adoption with revenue, while they figure out what sustainable unit economics might look like. The result is that providers either risk leaving money on the table by underpricing relative to value delivered, or bleed cash because heavy usage outstrips revenue.

Our [Global Software Study 2025](#) revealed two-thirds of companies that set revenue or profit targets for their AI initiatives achieved them. That sounds like success. But the fine print tells a different story: most of those targets were extremely modest, in the 0-10% range. Outcomes are incremental gains, not transformational ones. Just two years ago many executives were publicly forecasting that AI could drive 15-20% growth. The reality has fallen well short of those early expectations.



Our experience from talking about monetization with leading tech companies of the world they show just how urgent it is for the industry to move beyond guesswork and toward pricing frameworks that are grounded in outcomes, segmentation, and cost alignment.

AI has delivered some wins, but most companies are still under-ambitious, under-commercialized, or cautious. AI must move as a set of side projects delivering single-digit uplifts to a core driver of double-digit, business-wide growth.

How to drive profitable revenue growth through AI innovation

Driving significant revenue growth through AI innovation is not a given.

Think about GitHub's Copilot. It was a runaway success story: over 1.5 million users, huge adoption, and deep integration into developer workflows. It was literally writing half the code for its users, which is about as clear a product-market fit signal as you could imagine.

However, beneath that success, the economics didn't add up. GitHub priced Copilot like a classic SaaS product: \$10 per month, unlimited usage, predictable for the buyer. That makes sense if you think of software as a fixed-cost tool.

However, AI isn't fixed cost. Every query burns compute, every heavy user eats inference cycles, and some "super users" were consuming the equivalent of \$80 worth of compute for their \$10 subscription. On average, GitHub was losing more than \$20 per user.

Growth is supposed to improve margins. Here, growth made losses worse, and GitHub had many iterations on Copilot pricing to support adoption while making the model more sustainable over time.

Copilot's launch Model in June 2022

- \$10/month flat fee, unlimited usage.

Result: Heavy users consumed disproportionate resources, creating losses.

Early 2025

- **Pro Plan** with limitations on premium requests (baseline users pay less, but with guardrails).
- **Pro Plus at \$39** for power users, who are more compute intensive.
- **Business (\$19)** and **Enterprise (\$39)** tiers for organizations, providing scalable monetization aligned to team budgets and governance needs.
- **Simple auto-completion (low-cost):** Still accessible on lower tiers.
- **Premium requests (high-cost):** E.g., analyzing large codebases or context-heavy queries are fenced off to higher-paying tiers.

Copilot's model today

- **Copilot Free** – \$0;
2,000 completions + 50 chat messages/month, limited models and features.
- **Copilot Pro** – \$10/month or \$100/year (per individual)
Unlimited completions & chat with included models; 300 premium requests/month.
- **Copilot Pro+** – \$39/month or \$390/year
500 premium requests/month plus access to all available models and the coding agent.
- **Copilot Business** – \$19 per granted seat/month
Org management, policy controls, 300 premium requests/user/month.
- **Copilot Enterprise** – \$39 per granted seat/month
All Business features + GitHub.com-integrated chat, higher limits (1,000 premium requests/user/month), more enterprise controls.
- **Overages** – Additional premium requests on any paid plan are \$0.04 per request.

Copilot's updated model reinforces a segmentation strategy: light users get a generous free tier, while heavier or more specialized usage is increasingly fenced into higher-priced plans. Copilot Free stays accessible with limited monthly completions and chat, while Copilot Pro (\$10) unlocks unlimited baseline usage plus an allowance of premium requests. For the true power users, Copilot Pro+ (\$39) adds more premium volume, access to all available models, and the coding agent.

On the organizational side, Copilot Business (\$19/seat) introduces management and policy controls with a moderate premium-request budget, while Copilot Enterprise (\$39/seat) layers on deeper governance, GitHub.com-integrated chat, and significantly higher limits. Any plan that exceeds its premium-request allotment rolls into a clear overage price of \$0.04 per request.

Overall, GitHub continues to push a tiered framework: baseline completion and chat experiences remain affordable, while advanced capability and enterprise governance sit behind increasingly premium tiers. After a transition phase, higher consumption is now available across plans with usage-based pricing.

The Copilot case is a warning for every company rushing to launch AI features: adoption is not enough. If your pricing doesn't reflect the economics of AI-as-labor, scaling will hurt you rather than help you

Some claim that costs will eventually fall as infrastructure improves or models get cheaper. Others contend that each new generation of large language models becomes greedier. Both arguments miss the deeper issue: pricing models haven't evolved at the same pace as the technology. The problem isn't just cost, rather the lack of alignment between value delivered, cost incurred, and price charged.

AI monetization has to move away from flat, unlimited pricing. Tiered structures can be a first step to align cost with value and usage intensity. Light-touch users shouldn't subsidize heavy ones. Compute-heavy queries should be priced separately, the way enterprises already pay extra for advanced labor or premium expertise. Hybrid models, with a per user component and a consumption-based component, are gaining traction.

Clay Case Study:

Rapid shift to usage-based pricing and rapid revenue acceleration

Clay is another compelling illustration of how AI companies must evolve their pricing models as both the market and product usage mature. The company began with a straightforward user-based pricing structure: simple, fast to launch, and aligned with the early need to maximize adoption. But within less than two years, Clay had transitioned toward a pure usage-based model, reflecting a deeper understanding of how customers were actually deriving value.


This pricing evolution coincided with a pronounced acceleration in revenue. Clay grew from **\$2m in 2020 to \$7m in 2022** while iterating through hybrid models, then surged to **\$16m in 2023, \$31m in mid-2024, and \$37m by early 2025**. While not all of this growth can be attributed solely to pricing, the shift clearly enabled Clay to monetize high-intensity usage more effectively as AI adoption expanded.

Clay's journey showcases a broader truth for AI product builders: **your pricing model cannot remain static while your product, your users, and the maturity of AI evolve.**

What works at launch rarely remains optimal. Instead, the winners are those who continuously iterate, testing new structures, responding to usage patterns, and adapting as the market moves.

In short

Clay demonstrates that you must keep refining your pricing as the market, customer needs, and AI adoption evolve.



What is the best monetization strategy for your company?

Monetization can follow three main paths:

1.	Direct monetization with a price increase: AI is bundled into existing plans, but pricing is raised to reflect its added value. This works best when the AI is broadly relevant and variable costs are manageable.
2.	Direct monetization with a separate offering: AI is carved out as a standalone add-on, often tied to usage tiers or premium requests. This model is common when AI has high variable costs and usage intensity varies widely between customer segments.
3.	Indirect monetization without a price increase: AI is embedded to strengthen the core product, driving adoption, retention, and upsell of the broader platform. This approach makes sense when the AI capability is strategic but not cost-intensive. For example, features that enhance user experience or efficiency without significantly raising operating costs. AI drives indirect monetization through lower churn, higher cross-sell, and stronger conversion to premium tiers, even without an explicit price uplift.


The decision of how to translate adoption into sustainable revenue isn't one-size-fits-all.

It depends on three key considerations:

- **Ability to monetize indirectly:** Can the AI capability be folded into existing metrics (like seats, transactions, or API calls), or does it require its own billing unit?
- **Relevance of the AI capability:** Is the AI a broad, mass-market feature everyone will use, or is it a niche but highly valuable capability for a targeted audience?
- **Variable cost profile:** Does delivering the AI service carry high incremental costs (compute, inference, GPU time), or is it relatively low-cost at scale?

These factors shape which monetization method makes the most sense.

The important point is that all three models can succeed, but only if they are matched carefully to the product's economics and the customer's perception of value. Direct models make sense when costs and value are clear. Indirect models make sense when AI is better positioned as an enhancer of stickiness rather than a standalone revenue driver.




Our [Global Software Study 2025](#) shows that 45% of companies plan to use two or more monetization methods for AI. Instead of a universal playbook, we're seeing a portfolio of approaches that map to different product types, cost structures, and customer expectations.

- **Premium plan bundling (42%):** Equivalent of direct monetization with a price increase. It works when AI has broad appeal and relatively predictable costs, strengthening the upsell path.
- **Seat-based add-ons (41%):** The traditional SaaS meter. It's simple, intuitive, and easy for customers to understand, but it doesn't always match AI's value and variable cost profile.
- **Usage-based add-ons (39%):** This is where pricing begins to reflect AI's economics more closely. It ensures heavy users pay proportionally more, which protects margins.
- **Outcome-based add-ons (35%):** Pay for results, not features or usage.

No single method dominates. Each is being adopted by roughly a third to a half of companies. And that's not necessarily a problem. If anything, it highlights the value of experimentation. Especially in markets where pricing is not fully transparent, companies benefit from testing different structures to understand what drives the best traction, adoption, and monetization.

Monetization framework: Direct vs. indirect monetization



Source: Simon-Kucher

How to choose the right monetization approach

The monetization decision should ultimately align with two factors in our framework: expected customer adoption and business value delivered.

1. Roll-in with price increase	2. Create a premium tier:
<p>Use only when both conditions are strong:</p> <ul style="list-style-type: none"> AI functionality is widely needed, and it delivers high business value across the base. <p>This is the top-right quadrant: high adoption, high value. Even here, the risk is nontrivial: HubSpot initially pursued this path and later dialed it back, showing how sensitive customers can be to forced price increases.</p>	<p>Best when AI delivers high value, but adoption will not be universal. This reduces risk by monetizing depth of use rather than forcing breadth. ServiceNow and others have moved in this direction.</p>
3. Offer as an add-on (seat- or usage-based)	4. Roll-in with no price change
<p>A strong choice when:</p> <ul style="list-style-type: none"> value varies significantly across customer types, AI costs scale with usage or intensity, or you want a dedicated price metric that customers understand. 	<p>Only viable if:</p> <ul style="list-style-type: none"> AI is becoming table stakes, or marginal costs are low and predictable, or you already operate a usage-based backend that can absorb variability. <p>Zapier is an example of leaning toward table-stake AI rolled into existing plans.</p>

The interplay between expected adoption, business value, and cost intensity determines the right placement in the framework:

- **Roll-in with a price increase** → Only when extremely confident in universal appeal
+ high value
- **Premium tier or AI-specific tier** → Safer when value is high but adoption varies
- **Add-on (seat or usage)** → Ideal for variable value and cost intensity
- **Roll-in with no price change** → Appropriate only when AI is table stakes or costs are tightly controlled

As pricing transparency and usage patterns mature, companies benefit from testing multiple structures to understand where customers see the most value, and where monetization friction is lowest.

Case Study:

ServiceNow's clever monetization of GenAI drives double-digit growth

What they did	Rather than simply rolling AI into existing tiers for free, ServiceNow created differentiated packages (Pro Plus and Enterprise Plus) with embedded generative AI functionality. These AI-enabled packages reportedly carried about a 30% price premium compared to the base tiers.
Why it worked	<ul style="list-style-type: none"> ▪ Clear value creation: Customers experienced productivity improvements of 40–50%, making the ROI obvious and the willingness to pay very high. In their words: “There’s no price sensitivity around it because an improved productivity of 40–50% sells itself.” ▪ Fast adoption: AI features quickly became ServiceNow’s fastest-growing product launch ever. Adoption was nearly frictionless once the value was demonstrated. ▪ Shorter sales cycles: AI packages sold faster than comparable non-AI tiers, accelerating revenue capture.
The strategic evolution ahead	While tiered packaging has been effective, ServiceNow has already signaled its next step: moving toward usage-based pricing for AI agents. This reflects the broader market trend: early monetization often succeeds through bundling and tiering, but long-term sustainability will require models that align costs with usage and outcomes.

ServiceNow demonstrates the upside of incumbent-style AI monetization, enriching premium tiers to drive upsell and expansion. In contrast to GitHub Copilot’s cautionary tale of flat pricing and unit economics challenges, ServiceNow shows how premium-tier packaging, value communication, and customer willingness to pay can unlock transformational growth.

Incumbent vs. disruptor strategies

The “best” AI monetization strategy isn’t just about the AI capability, but also the company’s position in the market. Incumbents can afford to treat AI as an enhancement to stimulate migration into higher-value bundles. Disruptors must use AI itself as the differentiator. However, they also must be careful with unit economics, since they can’t cross-subsidize with other features the way incumbents can.



- **Incumbents:** Enrich existing premium tiers with AI, positioning it as a value booster that drives upsell within a broad package structure. Their challenge is not cannibalizing too much of their existing revenue while still making AI feel worth paying more for. At the same time, incumbents must distinguish between AI features that have quickly become table stakes (and therefore serve a defensive role) and those that genuinely drive incremental value. The former should reinforce core offerings, while the latter can be monetized within premium tiers to support meaningful revenue expansion.
- **Disruptors:** Fence AI capabilities directly, often through credits or usage caps, because they need tighter control of costs and don’t have as many levers for differentiation. Their challenge is avoiding commoditization while scaling.

Take the example of ServiceNow, which already has a deep feature stack. They don’t need to reinvent the wheel; instead, they weave AI into their premium packages to push customers up the value ladder. This works because incumbents can rely on breadth, using AI as one more feature in a rich package, reinforcing the “land-and-expand” motion.

On the disruptor side, companies like Cursor don’t have that same breadth of features to play with. Their differentiation comes more from how they fence usage and control costs. By packaging AI around credits, daily limits, or compute-heavy features, they can create clear price/value differentiation without needing a massive feature set. This is less about “land and expand” and more about “survive and scale”, controlling costs while signaling tiers of value to customers.

Pricing: More than just willingness to pay

In traditional SaaS, the lighthouse was always willingness to pay. Customers, even if they didn’t love the question, usually had a reference point. They knew what a CRM should cost, or what they paid for project management tools. Those frameworks worked because customers could compare value within familiar categories.

With AI agents, that reference point has shifted. They can’t say what they’d pay for “an autonomous agent that improves your customer communications,” because they don’t yet have a clear reference point for what that capability is worth or how consistently it delivers results. They don’t know how reliable the agent will be, how good the quality will

be, or how much of their workflow it can truly own. Willingness to pay needs to be re-grounded in fundamentals of value economics.

In the absence of clear customer benchmarks, companies need to lean more heavily on other inputs: economic value add, cost of alternatives, competitive positioning, and the realities of AI's variable costs. The task now is to help customers understand, quantify, and price that value.

This highlights a broader truth in pricing: no single input is enough. Instead, price levels should be a balance of different perspectives, ensuring prices are competitive, defensible, and reflective of the value delivered.

In the case of GenAI and agentic AI, the emphasis needs to shift more toward the right



- **Economic value add:** What's the actual saving in time, headcount, or efficiency for a given use case? Can the agent add a whole new revenue stream beyond savings? That may vary wildly by customer segment. An AI agent that saves an SMB a few hours a week is helpful; one that saves an enterprise 50 FTEs is transformational.
- **Cost of alternative solutions:** What's the true baseline the customer is comparing against? This includes both "do-it-yourself" options and external substitutes.
 - a) **Internal alternatives:** If the alternative is hiring and managing 20 analysts internally, that establishes a very different pricing reference than adding a low-cost plugin.
 - b) **External alternatives:** For many agentic use cases, the relevant benchmark increasingly goes beyond traditional software. Services priced outside the software category—such as BPO or managed services—become the dominant reference point. For example, if a call center charges \$1 per minute for customer calls, an


agent that performs the same task inherits that price level and model as a natural benchmark. In these cases, BPO pricing levels and models provide a powerful reference, allowing agentic AI providers to align with pricing structures customers already understand while demonstrating structurally lower costs driven by scale and automation

Of course, cost matters far more in AI than it ever did in SaaS. In traditional SaaS models, marginal costs were close to zero - each additional user was largely pure margin. In AI, by contrast, every inference consumes compute, and every complex query directly drives cost. According to a recent analysis by The Economist, OpenAI's demand for computing power - by far its largest cost - remains tightly coupled to the company's revenue. As a result, the cost floor is not an abstract consideration but an operationally critical constraint for AI solutions.

Meanwhile, competition and positioning still matter. You have to understand where you sit relative to peers, what your differentiators are, and what premium (or discount) you can command. But in this new market, you can't simply peg your price to competitors; you have to justify it through economics and alternatives.

Pricing is a moving target

Despite their new capabilities, most AI agents remain priced like traditional software, primarily on a per-seat or flat-fee basis, showing that the market's monetization logic hasn't yet caught up with the technology. Seat-based pricing makes up the largest share at 41%, followed by fixed fee packages at 30%. That means more than two-thirds of AI agents today are priced in ways that look very familiar to enterprise buyers: predictable, simple, easy to budget for, but not necessarily tied to the actual value or cost structure of the AI.



Usage-based pricing has traction at 23%, which makes sense for providers who want fairness (heavy users pay more) and need to protect themselves from the variable costs of inference and compute. But it's still a minority compared to seat-based or fixed-fee approaches.

Outcome-based pricing is only at 7%. It's the least common today, even though it is often the most effective at attracting customers. The appeal is clear: buyers see it as fair and low risk, paying only for measurable results. The challenge, however, is practical: outcome-based models are harder to implement, require strong attribution, and demand a higher level of customer trust.

“


Seat-based pricing remains widespread because most agents today are still in an **augmentation phase**. As agents become increasingly autonomous, however, seat-based models will begin to cannibalize themselves – ultimately becoming unsustainable for agent providers.

– Alexander Ammer, Partner

The implication here is that the market is in transition. Today's dominance of seat-based and fixed-fee models reflects several underlying dynamics:

1.	Customers are still thinking of AI as “software” which they’re used to buying per seat or per license.
2.	Vendors often prioritize adoption vs monetization in the first phase of introducing AI innovations – leading to unlimited usage as a first step.
3.	Many AI use cases are still augmentation rather than full replacement, making it harder to price directly on outcomes.
4.	Customers continue to seek predictability, something that usage-based or outcome-based models naturally make harder. This creates friction, especially for enterprises working with set budgets. The emerging answer is hybrid constructs: step-models, commit levels, or credit packs that provide usage flexibility while still giving buyers upfront visibility and control.

In today's world, companies can't assume the first pricing model they choose will be the one that lasts. Instead, they need to think of pricing as evolving from augmentation to replacement, and from fuzzy value to hard ROI. The more the AI capability migrates toward high impact, high tangible value, the more they need to orient around usage or outcomes, not seats.



AI monetization is not static. It has a trajectory. Companies may start with the comfort of SaaS-era models like seat-based pricing because it's familiar and easy to sell, but as AI agents mature, those models just won't capture the value.

Don't think of AI pricing as a one-time choice. Think of it as a migration path. Companies that adapt quickly will capture the upside. Companies that cling to the old models too long will leave money on the table or even burn cash. It's not just about what pricing model fits today but where the value delivery is going and how your monetization strategy can anticipate that shift.

Outcome-based pricing & packaging

Let's think back to the relic of boxed software in the 1980s–1990s. Originally, many businesses had the mindset that software could be treated as a one-off purchase: buy a perpetual license and “own it” forever.

Then, SaaS flipped that. No longer just a tool, it was something customers subscribed to, an ongoing service that kept delivering value. The subscription aligned better with the reality that software wasn't static, it needed maintenance, upgrades, and support.

Today, employers don't buy a piece of code once; they subscribe, continuously paying as they use it. And SaaS packaging has a few predictable levers:

- **Feature gating** → upsell by unlocking new capabilities.
- **Usage thresholds** → upsell by allowing more seats, more storage, more transactions.
- **Governance & security layers** → upsell to the enterprise tier where risk tolerance is lower and control is higher.

If software is a tool, then scaling usage is just scaling how many people get to hold the tool or how many times you can swing it. More seats, more storage, more transactions. Logical because usage has been the proxy for value.

Enterprise buyers also care about risk, compliance, and control, so adding that into the package was another way to upsell. It made sense to bundle those capabilities as extras because they weren't core “features” of the tool, but rather the guardrails around it.

Stepping into a new paradigm: Service as a Software.

If software companies were able to make the leap from allowing customers to “own” software outright (the boxed model) to “renting” it on subscription, why should the next step be unthinkable? What happens if we start viewing software as a worker that produces measurable outputs?

Nobody would say “I'm subscribing to Alex” or “I bought Jordan for a one-off fee.” Humans are valued through frameworks like:

- **How many hours they work (time-based)**
- **How much revenue they generate (performance-based)**
- **How scarce and irreplaceable their expertise is (premium pricing for specialists)**

Why should software be exempt from this logic, especially when it is mirroring or augmenting human tasks? Many companies today don't think in terms of, "I bought a thing with these features." They're thinking, "I need a job done, a problem solved, a service delivered, and AI software can do it." They're assessing value in terms of performance, reliability, and outcomes.

This feels intuitive, but it's radical compared to how most software pricing works today. Software companies still mostly charge by license, seat, or API call. That's like paying a human by the number of times they pick up the phone or how many words they say, instead of the outcome they deliver.

The closest analogy is no longer "tools and equipment." It's people. Because when you hire a human, you don't buy a feature list. You think:

- **Attributes:** Do they have the skills, personality, reliability, domain expertise, and potential I need? (Packaging)
- **Compensation model:** Do I pay them by the hour, by commission, or by salary? (Price meter)
- **Compensation level:** How much do I pay for their time or outcomes? (Price point)

Why shouldn't AI software follow the same logic? If an AI system is now acting like frontline support, a sales rep, or an expert analyst, then the way it is valued, priced, and packaged should mirror that too.

You don't subscribe to a person like you do your Netflix account. You compensate them for their contribution, their outcomes, and their scarcity.

“

For humans, we've had centuries of experience figuring this out. Employers have refined models of recruitment, retention, promotion, and compensation since the Industrial Revolution. We know how to package (job descriptions, qualifications), how to price (hourly vs. performance-based vs. expertise-based), and how to set levels (market rates, negotiation, seniority).

– Abde Tambawala, Partner

Case Study:

Intercom's outcomes-based pricing scales with value

The challenge	<p>Intercom wanted to monetize its new Fin AI Agent, a generative AI solution designed to reduce customer support volume by automating ticket resolutions. The challenge was to build trust in the product's effectiveness while creating a pricing model that scaled with adoption and value delivered. They also needed a model that could accommodate different customer profiles: those ready to go all-in on outcomes, and those who still relied on traditional seat-based pricing for parts of their support stack.</p>
The approach	<p>Instead of charging customers by seats, usage, or conversations, Intercom introduced an outcome-based pricing model:</p> <ul style="list-style-type: none"> ▪ Customers pay \$0.99 per successful resolution. ▪ Charges only apply when Fin AI fully resolves a support ticket. ▪ The model tightly aligns price with measurable customer value: Time saved, lower labor costs, and faster resolutions. <p>Crucially, Intercom did not arrive at this definition on the first attempt. They initially experimented with an “AI-confirmed resolution” metric, requiring the model to explicitly confirm a solved ticket. However, customers struggled to understand and trust this threshold. After testing and customer feedback, Intercom refined the definition to a more intuitive and widely accepted measure (confirmed or assumed resolution, depending on context). This iteration was key to landing on a model that felt both fair and value-aligned to buyers.</p> <p>After testing and customer feedback, Intercom refined the definition to a more intuitive and widely accepted measure: a dual definition of “confirmed resolution” (explicit positive customer confirmation) or “assumed resolution” (the customer leaves without further questions). This iteration was key to landing on a model that felt both fair and value-aligned to buyers.</p> <p>Importantly, outcome-based pricing is not applied in isolation. Intercom continues to monetize via seats where it makes sense, for example, human support agents using the platform, while allowing customers with fully AI-equipped workflows to choose a purely outcome-based model. This flexibility helps meet customers where they are on the AI adoption curve.</p>
The results	<ul style="list-style-type: none"> ▪ Rapid adoption: More than 17% of customers have already adopted outcome-based pricing for Fin. ▪ Clear ROI: At ~\$1 per AI resolution versus ~\$10 for a human resolution, companies see ~90% cost savings. Customers also report saving ~142 agent-hours per month.

The results

- **Scalable economics:** As usage grows, so does revenue. Fin's share of resolved tickets increased from 15% in the first month to 45% by month five, driving significant incremental revenue for Intercom.
- **Frictionless sales motion:** Because customers only pay for successful outcomes, procurement hurdles are lower, and it's easy for buyers to justify the investment internally.
- **Greater cross-company alignment:** Intercom found that tying revenue directly to Fin's performance rallied sales, product, and R&D teams around a singular metric (successful resolutions) reshaping internal priorities and accelerating product improvements.

The lesson

Intercom demonstrates that outcomes-based pricing can succeed at scale when outcomes are clearly measurable and tied to business impact. Intercom created a win-win model: customers capture efficiency gains, while Intercom monetizes in direct proportion to the value delivered.


Their journey also shows the importance of iteration: the right definition of success emerged only after testing with customers. And by keeping seat-based options where needed, Intercom preserved predictability for traditional teams while enabling AI-native buyers to scale purely on outcomes.

Pricing for levels of autonomy, impact, and sophistication

The term “AI agent” itself is still in flux. Sometimes it means a scripted bot, sometimes a copilot with light autonomy, sometimes a fully orchestrated system that can run across apps without supervision. The definitions haven’t yet settled.

In some ways that makes the moment even more powerful. Whenever categories are blurry, the race is wide open. Standards haven’t hardened. Business models haven’t solidified. And in that space, there’s an enormous opportunity.

Before deciding how to monetize, GenAI and agents should be evaluated along three dimensions: autonomy, attribution, and sophistication. These dimensions will help dictate which pricing approach is feasible and credible.



Agent autonomy

When we hire humans, we already think about what level of direction they need:

- At the **entry level**, you have junior analysts. You must break down tasks one by one, handhold, review their work, and then assign the next task. That's what today's narrow AI still feels like: point it at one specific task (summarize this, classify that, calculate this) and then the human manager has to stitch it all together.
- At the **mid-level**, you have experienced executors. You set a goal and outline the steps, and they execute reliably, checking in as they go. That's closer to current "workflow AI" or copilots: they need some structure, but they can still take a lot of the execution off your plate.
- At the **highest level**, you have leaders or directors. You tell them the goal, and they go off and figure it out, solving problems, coordinating others, and reporting back only when outcomes are achieved. That's like an "autonomous AI agent." It's the dream scenario where you say: Here's my business goal, go make it happen.

A similar logic can be applied to an AI agent. If the AI requires constant human prompting, you're still in augmentation mode, and user- or seat-based pricing might make more sense. But if the agent can operate with high autonomy, taking goals rather than instructions, then usage- or outcome-based pricing feels more natural, because the AI is acting more like an independent worker than a tool.

Attribution of impact

Capturing value depends on one critical question: can you attribute the outcome back to the agent?

Some AI agents directly drive measurable business results, like a sales agent closing deals or a customer support agent resolving tickets. The link between the agent's actions and revenue generation or cost reduction is clear, which makes outcome-based pricing possible. Others, such as a copywriter agent improving the quality of marketing content, influence outcomes more indirectly. Their contribution is valuable, but harder to quantify. Pricing for these agents often needs to fall back on proxies like usage, deliverables, or productivity metrics until attribution can be established.

This is critical for outcome-based models. If you can't reliably measure the agent's contribution, it's risky to charge on outcomes. You'd likely fall back on usage or activity-based metrics. But if attribution is straightforward and measurable, outcome-based pricing becomes both credible and compelling.

Sophistication


Domain expertise in humans is another clear way employers differentiate value. A generalist is useful, but interchangeable. You can find another analyst who "knows a bit of logistics." But when you bring in a domain expert, someone who's lived through the edge cases, understands the unspoken rules, and has scar tissue from real-world problem solving, their value multiplies.

If that expertise is in a rare, high-stakes niche (like ERP systems for hospitals, or regulatory strategy in biotech), the market rate for that person can skyrocket because there are so few people who can do it well.

Now, let's map this to AI. Some AI agents are generic and low in complexity; they don't warrant premium pricing. Others have high domain expertise, require specialized training, or consume heavy compute. Higher sophistication supports premium tiers, specialized packaging, or differentiated pricing, because customers recognize the depth and scarcity of the capability:

- **General-purpose LLMs** (like GPTs used off the shelf) are like generalist employees. They're flexible, can adapt to a wide range of tasks, and are cost-effective at scale. But they're not where the deepest value lies.
- **Domain-specialized AI systems**, trained or fine-tuned for a particular context, are the real equivalents of seasoned consultants. An AI that's been trained specifically for hospital ERP workflows, or tax compliance in a certain jurisdiction, becomes disproportionately valuable. It may not look as "broad" on a feature sheet, but to the buyer who needs exactly that, it's indispensable.


- **Rare domain AI experts.** Think of highly specialized vertical AI agents that integrate deep regulatory knowledge, niche industry data, or hyper-local compliance. These are scarce, mission-critical, and command a premium price point.



People assume AI will be cheaper than humans because it automates. But in many cases, the opposite should be true. AI isn't just a cheaper worker, it's a super-worker: faster, more precise, infinitely scalable. And when we look at AI pricing through this lens, it flips the usual assumption.

Agentic AI price metric spectrum

The agentic AI price metric spectrum can be understood as four distinct “flavors” of usage-based pricing, each anchored in a different mental model of value. On the far left, pricing is driven by the resources an agent consumes; on the far right, by the outcomes it delivers. Most providers fall somewhere between these poles, selecting metrics that match how their customers already think about value, how homogeneous their use cases are, and how predictable their costs need to be.



Resource-based pricing is the purest cost model and the one most familiar to hyperscalers and foundational AI providers such as AWS, Google, and Azure. Here, customers pay for compute time, tokens, or raw usage: simple meters that scale with consumption and align neatly with underlying cost drivers. While this model is highly transparent and easy to adopt across a wide range of workloads, it doesn't necessarily reflect the business value the agent creates. It works best when buyers are technically sophisticated or when the provider's differentiation lies in infrastructure rather than outcomes.

Moving one step closer to value is **activity-based pricing**, where the meter reflects the actions the agent performs: queries, credits, or workflow steps. This model retains the predictability of resource pricing but begins to connect usage to productivity. ChatGPT is a good illustration: while the company communicates "unlimited usage" at the base tier, the real monetization sits beneath the surface in the form of monthly limits for GPT-5 Pro or Deep Search queries. Business and Enterprise buyers who exceed those limits can purchase credit packs, essentially extending the activity meter. Other companies, like Picsart, push activity-based pricing much further, assigning explicit credit values to different AI actions (two credits for an AI filter, four for background editing, twenty for image-to-video generation). This approach allows for very precise monetization of innovation and cost protection, though it comes with added complexity that not all customers are willing to navigate.

The next progression toward value is **output-based pricing**, where customers pay for the work the agent produces rather than the work it performs. This model fits naturally in

sectors that already track outputs: voice agents priced per minute, or content-generation platforms like Synthesia or ElevenLabs charging per video or per audio asset. The challenge, however, is ensuring outputs are sufficiently homogeneous to support a stable price. Salesforce, for example, experimented with pricing its AI at €2 per conversation but reversed course within months. A “conversation” proved too variable: it might represent a simple support query or the close of a million-euro deal, creating both customer hesitation and internal challenges in sustaining the model. Salesforce ultimately shifted back to a credit-based system, reflecting the difficulty of anchoring price to heterogeneous outputs.

At the far right of the spectrum is **outcome-based pricing**, often framed as the holy grail because it ties revenue directly to measurable business impact: chargeback reduction, resolution rates, or other KPIs. Although highly appealing in theory and persuasive in new-logo conversations, outcome-based pricing is rarely straightforward in practice. Attribution must be airtight, value must be quantifiable, and customers must be willing and able to budget for a variable fee tied to results. Over time, these conditions can become increasingly complex to maintain, especially when multiple tools, teams, or processes contribute to the same outcome. As a result, outcome-based pricing remains more of an aspiration than a mainstream reality, reserved for use cases where value can truly be isolated and proven.

Across this spectrum, the right metric depends as much on industry norms as on product design. Resource and activity models thrive where buyers are already accustomed to technical meters; output models resonate in industries with established per-unit economics; and outcome-based models gain traction in markets that regularly outsource work or contract around performance. The most effective usage-based models are those that map cleanly to how customers already understand value, and to how reliably that value can be delivered and measured.

However, AI pricing isn’t just about what customers will tolerate. It’s about what your product can credibly support. If you overshoot (say, try outcome-based pricing when attribution is fuzzy), you’ll erode trust. If you undershoot (seat-based for a highly autonomous, high-value agent), you’ll leave money on the table. Choosing the right point on that spectrum depends on what kind of agent you’ve built.

Autonomy, attribution, and sophistication help companies map where their AI sits, and that positioning should drive their pricing model.



- Low autonomy + weak attribution + low sophistication? Probably stick to simple, SaaS-like metrics.
- High autonomy + strong attribution + high sophistication? You’re in the zone where outcome-based pricing can shine.

- High autonomy, low attribution = base fee + credit model on activities
- Low autonomy, low attribution, high sophistication = explore a higher fixed base, with a fair use policy on activities or outputs

Are customers ready for this shift?

Our latest Software Study suggests they are. 76% of AI buyers now prefer usage-based pricing, signaling strong appetite to move beyond traditional seat-based SaaS models toward models that reflect the actual value agents deliver. In other words, the market is ready to pay for what AI does, not just how many people use it.

Buyers' preferences mirror the same maturity path we've described:

- Early adopters still favor seat- and resource-based pricing, reflecting comfort with familiar SaaS logic.
- A growing share is moving toward activity- and output-based pricing, where usage and measurable outcomes drive value perception.
- And while only 10% of buyers currently prefer outcome-based pricing, this segment represents the leading edge: customers most comfortable linking price directly to business success.

This data reinforces the conclusion that the market's mental model is catching up to the technology. Providers who can operationalize usage- or outcome-based pricing will be ahead of both competitors and customer expectations.

How to adjust operations for agentic AI and GenAI

If outcome-based pricing is the “North Star” for AI agent monetization, and everyone agrees it’s where the market wants to head, then organizations need to start equipping themselves for the transition. Getting there is a lot harder than it sounds.

Any company considering outcome-based pricing needs to start by answering a few critical questions:

The first challenge is definition .	What exactly counts as an “outcome”? In customer support, is it a resolved ticket, or only if it avoids escalation? In sales, is it a closed deal, or a qualified lead? Without a definition customers recognize as fair, outcome-based pricing risks feeling arbitrary.
Then comes value .	Even if you agree on what the outcome is, what’s the right price for it? A CFO might ask: if one AI-resolved ticket saves \$10 of human labor, can we charge \$5? Or \$9? Or only \$1? Setting a consistent value anchor is tricky because perceived value varies by customer size, industry, and workflow.
The automation question is subtle but critical.	Can the agent carry a workflow to completion in a way that is measurable and attributable? If multiple agents or humans are involved, attribution gets murky. Without clear attribution, outcome-based pricing collapses back into usage-based because you can’t prove who delivered the result.
You cannot ignore buyer concerns .	Is our sales team able to position outcome-based pricing confidently, addressing customer concerns about budgeting and cost predictability? Buyers like the fairness of paying for outcomes, but they fear volatility in costs. If one month the AI resolves 5,000 tickets and the next month 50,000, the bill swings wildly. The CRO has to equip reps to sell this model confidently, overcoming concerns about budgeting and cost control.
Tracking outcomes at scale means building new metering systems.	Can we tie business results to product telemetry in a transparent, auditable way? That’s a big engineering ask.
Forecasting becomes exponentially harder.	Finance leaders will need new KPIs, scenario models, and investor narratives. Margins also fluctuate because compute costs may not scale linearly with outcomes. Can we forecast reliably when revenue is tied to outcomes we don’t fully control?

Everything we have discussed up to now all lands on one unavoidable point: this isn't a side project but a business model transformation that touches sales, finance, operations, product, and technology.

SaaS taught companies to think of software evolution as "just another feature wave." You add a module, launch a new integration, train sales, and keep going. But agents aren't just another feature. They blur the line between product and workforce, introducing variable costs, outcome-driven monetization, and new infrastructure.

For most incumbents, acknowledging this will mean admitting that existing systems such as billing, sales comp plans, forecasting models, and even cultural norms aren't fit for an agentic world. That's uncomfortable, because it means transformation, not incremental change.

Here, native AI disruptors have an advantage. They're building directly for this new reality, scaling globally from the start, embedding usage-based economics into their infrastructure, and designing billing and metering systems ready for continuous evolution. They don't need to retrofit legacy processes or mindsets. However, their challenge is different: they must prove credibility, earn trust, and commercialize sustainably before the market matures.

The result is an uneven playing field. Incumbents have scale, brand, and customer access but face transformation headwinds. Natives have agility and clean-sheet systems but must find defensibility before their innovations are copied or commoditized.

Considerations for the CRO

The predictability and repeatability of SaaS sales motions gave CROs a very stable foundation to run their revenue engines. When the mental model was "per seat, per tier," sales enablement was simpler: reps needed to know which feature set matched which buyer persona, forecasting was clean because ARR scaled linearly, and compensation could be tied neatly to new seat or license growth.

The very things that made SaaS sales so predictable are now the exact pain points CROs need to solve for in the AI era:

- **Sales enablement:** How do I arm my reps to explain usage-based or outcome-based models when customers aren't familiar with them? How do I help them size deals when usage can spike unpredictably?
- **Value communication & messaging:** Reps need to be trained to tell an outcome story: cost savings, hours saved, revenue uplift. This means the messaging toolkit has to evolve from "good, better, best" to "here's the ROI this agent delivers for your business."

- **Sales comp structure:** In SaaS, reps could be paid on predictable ARR targets. With agents, if revenue is variable (usage, outcomes), incentive plans must evolve to encourage both new logo acquisition and driving adoption/usage. CROs have to avoid plans that leave reps frustrated because customer usage, and therefore revenue, is outside their control.

Considerations for the CFO

For the CFO, the implications of GenAI and agentic models are profound because they strike at the heart of financial discipline: metrics, billing, and forecasting. The CFO can no longer rely on the old SaaS muscle memory. They'll need to manage the variability and align internal stakeholders (including the board and investors) on what "good" looks like in an agentic world.

- **KPIs and metrics:** The standard SaaS metrics no longer tell the full story. ARR and seat growth aren't enough when revenue comes from usage spikes or outcome-based fees. Finance leaders now need to monitor metrics like gross margin per inference, cost per resolution, or revenue tied to adoption intensity. They also need to create new investor narratives around healthy growth.
- **Billing systems:** Traditional systems built for predictable subscriptions struggle when asked to meter tokens, credits, or outcomes. CFOs must oversee the implementation of new billing infrastructure that can handle consumption-based pricing at scale, reconcile millions of micro-transactions, and ensure compliance with revenue recognition standards. Billing becomes a strategic capability rather than a back-office function.
- **Forecasting:** Perhaps the hardest shift is forecasting. With agentic AI, usage and outcomes fluctuate, making revenue streams volatile. CFOs now need more advanced forecasting tools, scenario planning, and closer alignment with product and customer success teams to predict adoption and usage patterns.

Considerations for the CTO

For the CTO, agentic AI is both a threat and an opportunity. It's a threat because the old SaaS tech stack won't cut it; it's an opportunity because those who solve metering, attribution, and scalable infrastructure will define the standards the rest of the industry follows.

- **Metering and tracking usage:** Outcome- or usage-based pricing requires precise, real-time metering. Product teams need to capture granular consumption data (credits, tokens, outputs, resolutions) across millions of micro-transactions, surface it transparently to customers, and reconcile it with billing. This is a major technical lift compared to traditional SaaS tracking, and if it isn't built well, it erodes trust.

- **Architectural changes:** Agentic AI workloads are compute-intensive and variable. Supporting them often requires new infrastructure choices, including GPU provisioning, orchestration of multiple models, data pipelines for fine-tuning. The architecture must also adapt to new pricing models (e.g., usage-based vs. outcome-based) so that the product can meter accurately, scale elastically, and remain cost-efficient. For many incumbents, this means re-engineering legacy SaaS architectures that weren't built for high-variability consumption.
- **Customer migration:** Moving from seat-based SaaS to agentic pricing is also a product and technical issue. Customers need visibility into their usage, predictability in their bills, and smooth transitions from old subscription plans to new agent-based models. That requires thoughtful design of dashboards, APIs, and migration paths that preserve trust and minimize disruption.

“

Outcome-based pricing is aspirational. It aligns perfectly with the value AI promises but it demands a level of organizational maturity that touches every function. It's less about setting a price and more about redesigning the company around measuring, attributing, and communicating outcomes

– Abde Tambawala, Partner

Turning into action

Over the course of this paper, we've seen how the rise of Agentic AI marks a fundamental break from the software models that came before it. Where traditional software acted as a tool and SaaS evolved into a service, AI agents are beginning to operate as intelligent, autonomous systems that execute work, make decisions, and deliver measurable outcomes. This shift redefines not only how software creates value but also how that value should be captured.

For companies, the implication is clear: the familiar SaaS playbook built around features, seats, and predictable renewals will no longer be enough. Agents blur the line between product and workforce, introducing new cost structures, new measures of performance, and new expectations for accountability and trust. The opportunity is immense, unlocking faster growth, higher productivity, and new forms of customer value. However, realizing it will require a transformation that touches every function, from product and pricing to finance, sales, and technology.

This brings us to the key questions for every leader navigating this transition:

- **What should companies do now to capture the promise of Agentic AI?**
- **And how can they adapt their business models before the market moves past them?**



A common formula for success:

Clarity of value + cross-functional alignment + the courage to move early.

This formula is what separates pioneers from followers in moments of business model transformation. Clarity of value is what makes pricing credible. It's also what builds trust internally. Once everyone in the company understands the economic impact of the AI, finance can forecast it, sales can sell it, and customers can justify it. Alignment is what turns AI from a collection of promising features into a repeatable, scalable business model. Finally, there's timing. The pioneers like ServiceNow, Intercom, and Salesforce didn't wait for the market to mature or for best practices to crystallize. They moved before perfect certainty.

The reality is that no one yet has the complete AI monetization playbook. Costs are still volatile, customer expectations are still forming, and models keep evolving. But waiting for stability means ceding the narrative to others. Courage doesn't mean recklessness; it means a willingness to test, iterate, and learn in public. The early movers don't have all the answers, but they are willing to find them faster.

Four imperatives that separate long-term winners from the rest

1. Redefine your value narrative.

The first step is to shift the conversation from what the technology does to what it delivers. Translating technical capability into clear business outcomes is the foundation of value-based pricing.

At Simon-Kucher, we help companies quantify and articulate that value, shaping the commercial story that drives both adoption and pricing power. We've supported hundreds of technology and AI clients in redesigning value propositions, typically achieving 20-30% revenue uplift through clearer value communication.

2. Design new pricing and packaging architectures.

The second imperative is to rebuild the monetization model for the AI age. This requires customer insight, testing, and careful calibration of economics and experience.

Across industries, Simon-Kucher has pioneered business model innovation, from bringing subscriptions to media and mobility to introducing outcome-based pricing in software and services. We help clients build, test, and implement segmentation, pricing model design, and monetization pilots that validate real-world performance.

3. Operationalize across functions.

AI monetization cannot succeed in silos. CRO, CFO, and CTO teams must align around shared metrics, compatible systems, and common definitions of value. That means retraining sales to sell outcomes, enabling finance to forecast variable revenue, and equipping product and technology teams to measure usage and outcomes accurately.

Our cross-functional programs help organizations adapt sales enablement, metrics, and infrastructure to ensure sustainable growth and profitability. We work hand in hand with our clients to not only design the best strategy but also help bring them to market.

4. Future-proof your model.

Finally, companies must design their monetization approach to evolve with the market. Commoditization, cannibalization, and cost volatility are inevitable in any fast-moving technology cycle. The challenge is to anticipate these pressures before they erode margin.

We help clients scenario-plan, benchmark against peers, and build resilient monetization strategies that protect profitability as the AI landscape matures.

Translating technical leaps into commercial innovation

For decades, software innovation has been defined by technical leaps. The next defining advantage will come from commercial innovation: how effectively companies translate AI capability into sustainable, profitable business models.

“

Technology defines what's possible.
Monetization decides what's profitable.

– [Alix Nepveux, Partner](#)

At Simon-Kucher, we believe that great technology deserves great monetization. Our role is to help companies make that connection, bringing together deep pricing expertise, cross-functional alignment, and decades of experience turning innovation into measurable growth. We guide clients through the complexity of this transition: clarifying value, designing next-generation pricing models, and embedding them across sales, finance, and product to ensure they scale.

Authors

**Alix Nepveux**

Partner

Office Paris

Mobile: +33 6 09 95 03 67

alix.nepveux@simon-kucher.com

**Abde Tambawala**

Partner

Office San Francisco

Tel.: +1 650 713 9372

abde.tambawala@simon-kucher.com

**Alexander Ammer**

Partner

Office Munich

Tel.: +49 89 544793 680

Mobile: +49 151 17646541

alexander.ammer@simon-kucher.com

Paris office

3 Square Edouard VII
75009 Paris, France
Tel. +33 1 566923 90
paris@simon-kucher.com

San Francisco office

100 1st Street, Suite 2200,
San Francisco, CA 94105,
United States
Tel. +1 650 641 4300
san_francisco-office@simon-kucher.com

Munich office

Luise-Ullrich-Str. 14
80636 Munich, Germany
Tel. +49 89 544793 0
munich@simon-kucher.com


Follow us on



Discover more on

simon-kucher.com

© 2025 Simon-Kucher & Partners Strategy & Marketing
Consultants GmbH. All rights reserved.
Photo credits: Gettyimages, istock, Shutterstock



SIMON 
KUCHER
Unlocking better growth