Selling work, not software: the AI business models start-ups are building

AI start-ups recombine classic patterns – what shifts is the billing unit, from seats to tasks and outcomes

Falk Uebernickel
University of St.Gallen (HSG)
ZHAW School of Management and Law

Human thinking, AI-assisted writing: the ideas and the argument are the author's; the text was written together with AI.

Abstract. Which business models does artificial intelligence (AI) create, and which are old models in new clothes? This note summarises what we found when we coded the self-descriptions of 6,347 Y Combinator (YC) start-ups from 2005 to 2026 against an established catalogue of business-model patterns. Three findings stand out. First, the launch of ChatGPT marks a break: the share of AI-native start-ups rose from 35% of the 2022 batches to 61% in 2023 and 82% in 2026. Second, the pattern mix moved from marketplaces to integrators. Third, we propose ten candidate AI business-model patterns – led by Agent-as-a-Service and Service-as-Software – that describe how AI firms sell work rather than tools. Our position is that almost all of these patterns recombine classic ones; what is genuinely new is the unit in which value is billed. We set out the data, the patterns, a twelve-model view by where value is captured, implications for incumbents and for teaching, and the limits of the evidence. The candidate patterns are hypotheses, not established findings.

1 Why look at start-ups

Business model innovation is mostly the recombination of a limited set of recurring patterns (Gassmann et al., 2014); a business model describes how a firm creates, delivers and captures value (Teece, 2010). New technologies rarely invent new patterns from scratch. They change which combinations pay off. If AI does that, the change should show first where new firms are designed on a blank sheet: in start-up accelerators.

Y Combinator is a useful early indicator. It has run two or more batches a year since 2005, and its public directory describes each firm in a single pitch-like line.[1] The directory is not a representative sample of all new ventures, but it is a long, consistent series of how founders frame their businesses. Figure 1 shows the frame shifting. AI-native firms (AI is the core of the product) made up 14% of the 2015 batches and 22% in 2020. With the first batch after ChatGPT (W23) the share jumped to 61% and reached 82% in 2026. On a logit scale the jump is +1.09 (95% CI 0.84–1.34) on top of the trend, and it occurs in every sector.

2 How we found the patterns

Data and coding. We took a snapshot of the YC directory on 31 October 2026 (6,358 entries; 6,347 with a batch year 2005–2026). A local language model, run on our own hardware, coded each self-description: AI status, industry, customer type, revenue mechanics and one to three classic patterns from the St.Gallen Business Model Navigator (Gassmann et al., 2014). A second pass added pricing logic, AI depth, the function automated by AI and a flag for firms whose logic the model did not see covered by any catalogue pattern.

[1] Data: public Y Combinator company directory (via the open yc-oss mirror), own analysis; the underlying data are not redistributed. Source: Y Combinator (2026); classification of the firm self-descriptions with a local language model (gemma4:26b); within-batch-year comparisons; no causal claims. The data are not to be redistributed.

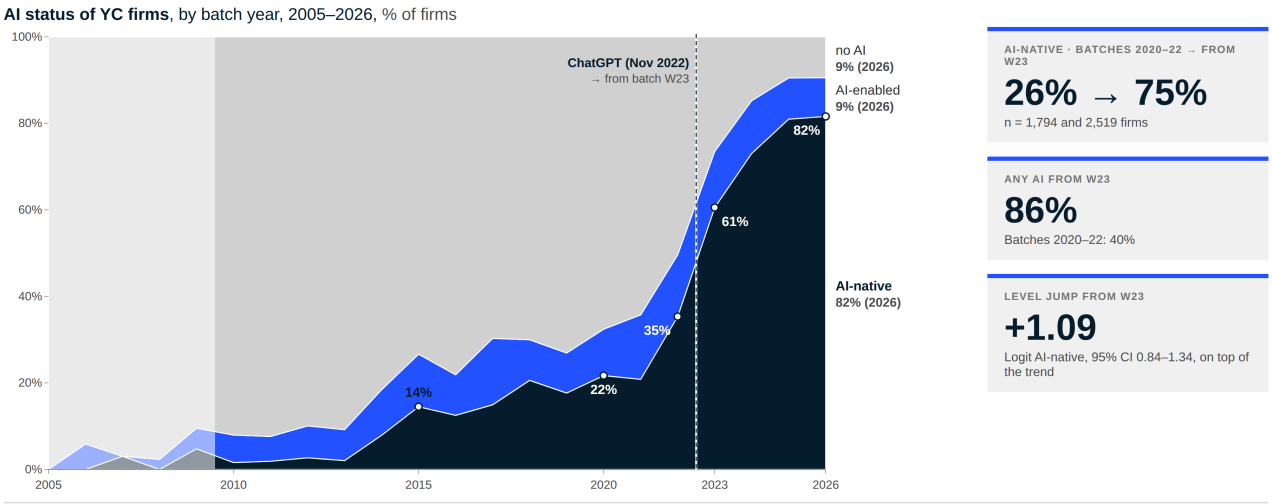


Figure 1. The post-ChatGPT break: AI status of YC firms by batch year, 2005–2026, in % of firms (n = 6,347). AI-native = AI is the core of the product; AI-enabled = AI is a relevant feature. Model coding of today's self-description (gemma4:26b, codebook v1.1); QC error of the AI flag 7.5% (only "AI-enabled" affected). Older firms have partly rewritten their profiles towards AI, so early years tend to be overstated; years before 2010 have fewer than 50 firms per year. No causal claim. Data: public Y Combinator company directory (via the open yc-oss mirror), own analysis; the underlying data are not redistributed.

Hand checks of random samples gave error rates of 9% for a firm's characteristic pattern, 7.5% for the AI flag and 0% for customer type and revenue mechanics.

Where the candidate patterns come from. The candidates did not emerge from the data on their own; the author defined them as hypotheses and then tested them on the data. The procedure had three steps. *Screening:* the second coding pass flagged 582 firms (9%) whose logic the model did not see covered by any catalogue pattern (hand-check error of the flag 3%). We also let the model name the uncovered logic in free text and clustered these names, but they collapsed onto the two examples given in the prompt ("agent-as-a-service" 255, "service-as-software" 211) and the clusters were barely separated (silhouette 0.14). An inductive derivation from the data was therefore not possible. *Definition:* the author then set ten candidate patterns, each with a one-line definition and the classic patterns it builds on. *Test:* the model assigned the candidates to the 582 flagged firms from this closed list (several labels allowed), with an open category "other". A hand check of 30 firms found 28 correct, an error of 7% (95% CI 2–21%). Only ten firms landed in "other", and a clustering of their texts showed no further frequent pattern. The data thus show that the ten candidates cover the flagged firms almost completely, how often each occurs and since when; they do not show that this is the only possible cut. Because the two leading candidates were also the examples in the screening prompt, their lead should be read with this priming in mind. The ten candidates are the author's framework.[2]

3 The AI business-model patterns

Table 1 lists the ten candidates with their logic, the classic patterns they build on and the number of YC firms carrying each label. Two dominate. *Agent-as-a-Service* (at least 406 firms, 15.5% of all AI-native YC firms) deploys an AI agent in a role – sales development, dispatch, claims handling – and positions it like staff. *Service-as-Software* (at least 144) sells the finished service, for example tax, accounting or immigration law, delivered by AI with a few professionals. The two often appear together (73 firms), and *Human-in-the-loop* services, where experts check and take liability for AI output, attach to both. The candidates are young: for seven of the ten candidates, the first quarter of firms arrived from 2023 on,

[2] Framework: Falk Uebernickel (candidate AI business-model patterns, derived from YC data). Counts are lower bounds, because only the 582 flagged firms were labelled.

Candidate pattern	Logic in one line	Builds on (classic)	YC n	Example (batch)
Agent-as-a-Service	An AI agent takes over a role or workflow and is deployed like staff	Subscription + Pay Per Use; Solution Provider	≥406	Floworks (W23)
Service-as-Software	The firm sells the finished service, delivered by AI with few experts	Solution Provider + Integrator	≥144	Caribou (W19)
AI infra for developers	Tools for firms that build AI or agents (payments, memory, evals, data)	Layer Player + Pay Per Use	≥68	Doublezero (S24)
Human-in-the-loop service	AI does most of the work, experts check and are liable	Solution Provider + Layer Player	≥50	superglue (W25)
Data flywheel	Use generates proprietary data that keeps improving the product	Leverage Customer Data + Lock-in	≥38	Exin Therapeutics (W25)
Outcome-based pricing	Paid per completed task or result; provider carries execution risk	Performance-based Contracting	≥29	GitStart (S19)
Autonomous ops	Robots, drones or autonomous systems are operated; the operation is sold	Performance-based Contracting + Rent Instead of Buy	≥28	InLoop Robotics (X26)
AI-native marketplace	AI or agents do the matching and negotiation	Two-sided Market + Orchestrator	≥17	FleetWorks (S23)
AI content as product	AI-generated content or products are what is bought	Mass Customization + Digitization	≥16	Promi (S24)
Model as product	The own model is the product, billed per token or call	Pay Per Use + License	≥3	Kailash Labs (F26)

Table 1. Ten candidate AI business-model patterns. Framework: Falk Uebernickel (candidate AI business-model patterns, derived from YC data). n = YC firms carrying the label among the 582 flagged firms (lower bounds; model providers billing per use were mostly coded as classic Pay Per Use and are missing). Batches: W = winter, S = summer, X = spring, F = fall. Examples are single self-descriptions, not checked case studies. Label error 7% (95% CI 2–21%). Data: public Y Combinator company directory (via the open yc-oss mirror), own analysis; the underlying data are not redistributed.

and 85% of Agent-as-a-Service firms come from batch W23 or later.

4 Twelve AI business models by where value is captured

The YC candidates describe start-ups. For the lecture we combine them with dated launches by established and well-funded firms into twelve AI business models, grouped by who pays. *Infrastructure and models:* builders pay for compute and tokens – AI infrastructure (e.g. Amazon SageMaker, 2017), the model API billed per token (OpenAI API, 2020, OpenAI, 2020), open weights with paid services (Llama 2, 2023) and model marketplaces (Azure AI model catalogue, 2023). *Applications and work:* firms pay per seat, per task or per outcome – the copilot seat (GitHub Copilot at USD 10 per user and month, GitHub, 2022), pay per outcome (Intercom Fin at USD 0.99 per resolution, Intercom, 2023) and the digital worker or Service-as-Software (Salesforce markets Agentforce as "digital labor", Salesforce, 2025). *Distribution and monetisation:* third parties pay – ads in AI answers (Google AI Overviews, 2024), agentic commerce (Perplexity Buy with Pro, 2024) and training-data licensing (publisher deals with AI labs from 2023). *Physical world:* AI is sold as a ride or inside a device – autonomous operations (Waymo One, 2020) and AI bundled into the device (Samsung Galaxy AI, 2024). All twelve build on classic patterns; only agentic commerce and the digital worker are still emerging. The infrastructure model corresponds to the "AI development facilitator" archetype of Weber et al. (2022). Figure 2 shows the twelve models side by side.

5 What shifts – and what does not

The billing unit. Our central position is that AI changes less *which* patterns firms use than *what* they bill for: from seats and licences towards tokens, tasks and outcomes. The YC data support this, cautiously. Among all firms in the 2025–26 batches, 4.2% of AI-native firms

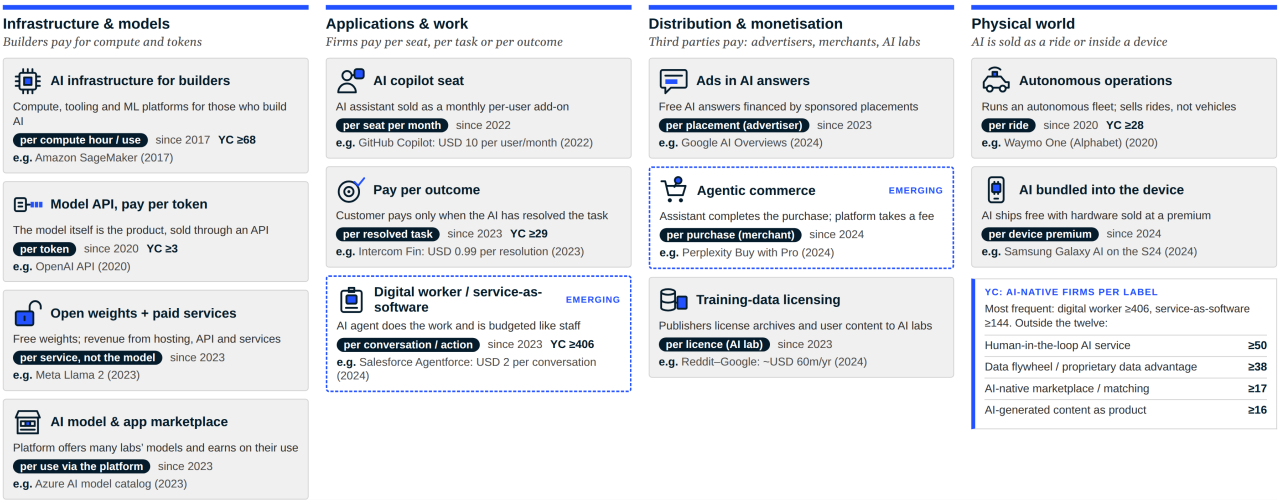


Figure 2. Twelve AI business models, grouped by where value is captured (schematic). Each tile: logic, billing unit, first dated launch, one example; dashed = emerging, not yet in a pattern catalogue. Framework: Falk Uebernickel (candidate AI business-model patterns, derived from YC data). Prices and dates from vendor pages and press; YC $\geq n$ = AI-native YC firms whose self-description carries the label (lower bounds; labels are hypotheses, hand-check error 7 %, 95 % CI 2–21 %). Data: public Y Combinator company directory (via the open yc-oss mirror), own analysis; the underlying data are not redistributed.

name payment per outcome as their main pricing logic (non-AI: 0%) and 5.6% usage-based pricing (non-AI: 1.5%). Subscription remains the default, and only few firms state their prices at all. Listed firms are slower still: in 10-K filings, outcome-based language stays around 1% and is mostly asset-manager fees and pharma (U.S. Securities and Exchange Commission, 2026).

Integration over marketplaces. Figure 3 shows the broader pattern shift. Between 2012 and 2026 (fitted trend), the share of YC firms carrying the Integrator pattern rose from 8% to 43%, Layer Player from 14% to 31% and Pay Per Use from 5% to 22%; Two-sided Market fell from 28% to 4%, Digitization from 23% to 6%. Start-ups are again building whole value chains instead of matching others' supply and demand. The "AI-native marketplace" candidate counts only 17 firms. Controlling for the number of patterns per firm, only one coupling is genuinely more frequent in AI firms: Integrator + Pay Per Use (observed/expected 1.23 versus 0.78 in non-AI firms). The often-cited Integrator + Solution Provider pair is in fact tighter in non-AI firms; its raw rise comes from the Integrator boom, not from AI.

Selling work instead of tools. Agent-as-a-Service, Service-as-Software and Human-in-the-loop services share one logic: the customer buys work done, not software to do the work. This is consistent with the operating-model argument that AI shifts scale from people to algorithms and data (Iansiti and Lakhani, 2020) and with conceptual work on generative AI as a driver of business model innovation (Kanbach et al., 2024). Note, however, that "priced like staff" is rarely stated in the self-descriptions.

What does not shift. Two popular claims find no support in our data. AI teams are not smaller: within the same batch year, AI-native firms have 0.97 times the headcount of non-AI firms, and tiny teams are typical of all young start-ups. And recombination stays largely classic: AI-native firms are not more often atypical recombiners than others. Every firm is coded with existing catalogue patterns, so some continuity is built into the method.

6 Implications for incumbents and for teaching

Incumbents. Start-ups have become AI-native faster than listed firms describe AI as part of their business. In 2026, 40% of 10-K filings name AI only as a risk factor. Among 5,701 listed

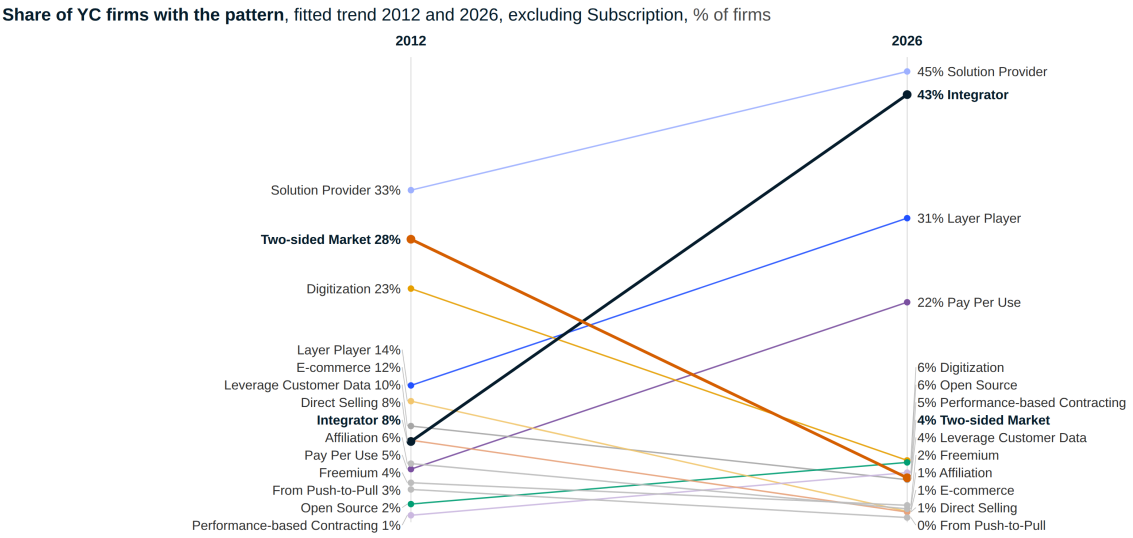


Figure 3. From marketplace to Integrator: share of YC firms carrying a pattern among their one to three patterns, fitted logit trend 2012 and 2026, in % of firms (n = 6,347; patterns with a significant trend and a change of at least 2 percentage points; Subscription, set for SaaS by codebook rule, omitted). QC error: correct pattern among the one to three 3%, rank-1 pattern 21%. Early batch years incomplete (survivorship); no causal claim. Data: public Y Combinator company directory (via the open yc-oss mirror), own analysis; the underlying data are not redistributed.

firms, only 1.2% (95% CI 1.0–1.5%) show a sustained switch to recurring-revenue language between 2018 and 2026, after removing de-SPAC and shell artefacts (U.S. Securities and Exchange Commission, 2026). As a hypothesis for discussion: the threat to incumbents is less a new pattern than a new billing unit – a competitor who charges per resolved case or completed task makes per-seat revenue look expensive, and moving a revenue model is rare and slow.

Teaching. For executives with little prior knowledge, three points follow. First, the classic pattern catalogues remain a useful toolbox (Gassmann et al., 2014): AI business models can be explained as combinations of known patterns – the new part is how they bill. Second, the design question moves to the value-capture side: what is the unit of value – a seat, a token, a task, an outcome – and who carries the execution risk? Third, the candidate patterns are a structured hypothesis to test on cases, not a catalogue to copy.

7 Limitations

The evidence has clear limits. YC firms are a selected, US-centred sample, not "all new business models". Everything rests on today's self-descriptions: pitch language, rewritten after pivots and maintained by surviving firms, which is why AI terms appear even in early batches and why a candidate's "first batch" can be too early. Coding was done by a local language model with measured but non-zero error rates, and the candidate labels have a hand-check error of 7% (95% CI 2–21%) on a small sample. The candidate patterns were set by the author and tested, not derived inductively. Counts are lower bounds. Re-running the classifier changes about 8% of rank-1 patterns. All comparisons are descriptive and within batch years; none is causal. The 10-K measures capture language, not revenue or use, and 2026 is a partial year. The twelve-model view combines the YC candidates with dated public examples; those examples illustrate the models and are not a systematic sample.

References

Gassmann, O., Frankenberger, K., and Csik, M. (2014). *The Business Model Navigator: 55 Models That Will Revolutionise Your Business*. Pearson / FT Publishing, Harlow.

GitHub (2022). GitHub Copilot is generally available to all developers. https://github.blog/news-insights/product-news/github-copilot-is-generally-available-to-all-developers/. GitHub Blog, 21 June 2022.

Iansiti, M. and Lakhani, K. R. (2020). *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*.

Harvard Business Review Press, Boston, MA.
Intercom (2023). Fin resolutions. https://www.intercom.com/help/en/articles/8205718-fin-resolutions. Help centre article.
Kanbach, D. K., Heiduk, L., Blueher, G., Schreiter, M., and Lahmann, A. (2024). The GenAI is out of the bottle: generative artificial intelligence from a business model innovation perspective. *Review of Managerial Science*, 18(4):1189–1220.
OpenAI (2020). OpenAI API. https://openai.com/index/openai-api/. Announcement, 11 June 2020.
Salesforce (2025). Salesforce introduces new flexible Agentforce pricing to accelerate the digital labor revolution. https://www.businesswire.com/news/home/20250515332990/en/Salesforce-Introduces-New-Flexible-Agentforce-Pricing-to-Accelerate-the-Digital-Labor-Revolution. Press release, Business Wire, 15 May 2025.
Teece, D. J. (2010). Business models, business strategy and innovation. *Long Range Planning*, 43(2–3):172–194.
U.S. Securities and Exchange Commission (2026). EDGAR full-text filings, Form 10-K, 2018–2026. https://www.sec.gov/edgar. Texts via KL3M Data: SEC EDGAR Form 10-K Filings, v1.2, ALEA Institute, CC BY 4.0, https://huggingface.co/datasets/alea-institute/kl3m-data-edgar-10-k.
Weber, M., Beutter, M., Weking, J., Böhm, M., and Krcmar, H. (2022). AI startup business models: Key characteristics and directions for entrepreneurship research. *Business & Information Systems Engineering*, 64(1):91–109.
Y Combinator (2026). Y Combinator company directory. https://www.ycombinator.com/companies. Public company directory; snapshot of 31 October 2026 via the open mirror https://github.com/yc-oss/api, own analysis.