HISTORICAL STUDY

When Does a Service Become a Product?

Most companies think they have productized a service when they start selling packages. The historical record points somewhere else, and it is less about standardizing the work than about what the provider can commit to before the work begins.

PUBLISHED
August 2026

CATEGORY
Operating models, service economics

EVIDENCE BASE
Filings, regulatory records, contemporaneous reporting

An agency sells four articles a month. Fixed scope, fixed fee, fixed cadence. On the invoice it looks like a product.

Underneath, a senior strategist is still deciding, every month, for every client, whether four articles are the right thing to be doing at all. Perhaps the site has a technical problem that no amount of publishing will fix. Perhaps the authority gap is the constraint. Perhaps the correct answer this quarter is nothing.

If that decision is being made fresh each time by an expensive person, the agency has a product-shaped price sitting on top of a service-shaped cost structure. It will feel scalable right up until it does not.

This pattern is not confined to marketing. It shows up in bookkeeping, legal work, recruiting, software implementation and installed home services. And the failures it produces are unusually well documented, because companies that price like a product and cost like a service tend to discover the mismatch in public.

THE DISTINCTION THAT RUNS THROUGH THIS STUDY

Standardizing what you deliver is not the same as standardizing what it costs you. A great deal of what gets called productization is the first thing described as though it were the second.

The question worth asking is narrower than it first appears. It is not whether a service can be made repeatable. It is whether the provider can commit to the economics of a transaction before knowing exactly what that transaction will require, and can do so while keeping expensive human judgment to a small share of the total work.

> Packages standardize what the customer buys. Productization standardizes what the provider can economically promise.

That is the recurring constraint. Not whether the work can be standardized, but whether the company can commit to the economics before it knows what the job will require. History suggests five ways of getting there. Only one requires making every individual job predictable, and it is the one everybody tries first.

THE FIVE ROUTES

Observe	See the variable that decides cost before committing to a price.
Control	Own the interface through which uncertainty reaches the work.
Pool	Let individual cost vary while cohort cost stays predictable.
Exclude	Publish what the standard path will not handle.
Price the outcome	Charge for successes when a failed attempt costs almost nothing.

EXAMINED IN FULL BELOW, AFTER THE EVIDENCE THAT PRODUCED THEM

THE SHAPE OF THE PROBLEM

Four states, not two

Businesses are usually described as either services or products. In practice there are intermediate states, and most companies that believe they are in the third are in the second.

01	02	03	04
Bespoke	**Packaged**	**Productized**	**Product-like layer**
Every engagement is diagnosed individually. Price follows effort, or follows a guess. Senior time grows with revenue.	Deliverables and prices are fixed. The decision selecting those deliverables is still made from scratch for each customer.	The provider can commit to the economics before the job starts, through observation, control, pooling, exclusion or outcome pricing.	The service stays bespoke. A layer inside it becomes the product, and captures the economics from a different position.
COST UNKNOWN AT SALE	THE COMMON MISTAKE	ECONOMICS COMMITTED	OFTEN THE BETTER MOVE

The gap between the second and third states is where money is lost. A packaged service carries the pricing risk of a product and the cost behavior of a service, which is a combination that works in good months and compounds badly in bad ones.

CASE

Two ways to price something you cannot predict

Zillow and Frontdoor both committed to prices without knowing what an individual transaction would cost. One of them lost roughly $881 million in its Homes segment in 2021 and closed the business. The difference is instructive, and it is not about the quality of the algorithm.

ZILLOW OFFERS, 2018 TO 2021

Zillow made instant cash offers on homes, typically within about two days, priced from a national model. The cost drivers that determine whether such an offer is profitable are local: street-level resale value, renovation cost, neighborhood demand, the direction of the market in that submarket over the following months.

In the third quarter of 2021 Zillow bought more houses than in the previous eighteen months combined, and continued buying into a decelerating market. In November it announced the closure of the business. Around a quarter of the workforce was cut, and some seven thousand houses were left to dispose of. The chief executive attributed the exit to the unpredictability of forecasting home prices being far greater than expected.

Read structurally, the problem was not that a model made errors. Models make errors. The problem was that each error was locked into an asset held for months, with no mechanism to correct the price after the fact.

$881M

Zillow Homes-segment loss, 2021, the year the business closed

COMPANY DISCLOSURE

~7,000

Homes remaining to be disposed of at wind-down

COMPANY DISCLOSURE

FRONTDOOR, HOME WARRANTIES

Now consider a business with the same structural problem and none of the consequences.

Frontdoor sells home warranty contracts under American Home Shield and related brands. It has no idea which member will call next month, or whether the failure will be a dishwasher or a compressor. Individual cost is genuinely unknowable, and always will be.

In 2025 the company reported revenue of about $2.09 billion, gross margin of about 55%, an adjusted EBITDA margin of roughly 26%, and around 76% of revenue coming from renewals. Its finance chief has attributed margin movement to the mix of price realization, the number of service requests per member, and cost inflation. That middle item is the one worth pausing on. Frontdoor tracks how often members call as an operating variable, because in aggregate it is stable enough to plan around.

$2.09B

Frontdoor revenue, 2025

COMPANY DISCLOSURE

55%

Gross margin on a cost base it cannot predict per customer

COMPANY DISCLOSURE

76%

Share of revenue from renewals

COMPANY DISCLOSURE

> A company does not always need to know what one customer will cost. Sometimes it only needs to know what a hundred thousand customers will cost.

This is pooling, and it belongs on the list of routes to product-like economics alongside the more familiar ones. It is not available to everyone. The historical cases where it works share four features.

The pool has to be large enough that incidence stops being noise. The contract has to reset often enough that a pricing error costs one cycle rather than one asset, which is precisely what Zillow lacked. Adverse selection has to be constrained somehow, through exclusions, service fees at the point of claim, eligibility rules or channel mix. And the provider needs a reliable way to observe how the pool is behaving, which is why claims accrual sits among Frontdoor's disclosed critical accounting estimates.

Take away the ability to reprice and pooling becomes a bet. Zillow held a portfolio too.

CASE

Where the judgment goes

The intuitive account of productization is that expertise gets removed. Across the historical record, expertise gets relocated, and knowing where it lands is more useful than pretending it left.

FANNIE MAE DESKTOP UNDERWRITER, FROM 1995

Before Desktop Underwriter, mortgage underwriters interpreted centrally written lending guidelines, one file at a time, thousands of times a day, with the inconsistency that implies. Fannie Mae put the guidelines into an expert system and moved it into production in the mid-1990s.

The system is often described as having automated underwriting. What actually happened is more specific and more useful. Fannie Mae's own operating guidance still instructs lenders to exercise prudent underwriting judgment, to confirm the accuracy of the data they submit, and to identify derogatory or contradictory information that the system did not analyze. Files the system refers rather than approves go to a human.

So judgment did not leave. It stopped being exercised independently on every routine case and concentrated into four places: the policy group that writes and reissues the rules, the data quality work at the point of submission, the design of the system itself, and the exception queue.

> The expert stopped deciding each case and started deciding the class of cases, once.
>
> THE PATTERN THAT RECURS IN EVERY SUCCESSFUL CASE IN THIS STUDY

ALIGN TECHNOLOGY AND ARAVIND EYE CARE

There is a second correction that matters more for professional services.

Align Technology manufactures Invisalign aligners. Its filings state plainly that because treatment is prescribed by a doctor, no two cases are alike, and that the doctor retains control over the treatment plan and participates in the customized design. That is irreducibly bespoke clinical judgment. The company serves roughly 291,000 doctor customers and behaves, economically, like a manufacturer.

Nothing about the doctor's decision was standardized. What changed is how much of the total transaction that decision occupies. Align has spent two decades compressing it, most recently with a treatment planning feature announced in

October 2025 that generates a doctor-ready initial plan in about fifteen minutes rather than days. The doctor still approves, and the approval still gates manufacturing.

Aravind Eye Care reaches the same place from a different direction. It did not simplify the surgeon's decisions. It removed from the surgeon everything that was not a decision: eye washing, anesthetic injection, patient preparation, refraction, preliminary tests, all reassigned to paramedical staff working alongside two operating tables. Reported surgical output moved from roughly five or six procedures a day to twenty-five or thirty, at complication rates reported at half those of the UK National Health Service.

ILLUSTRATIVE, NOT MEASURED

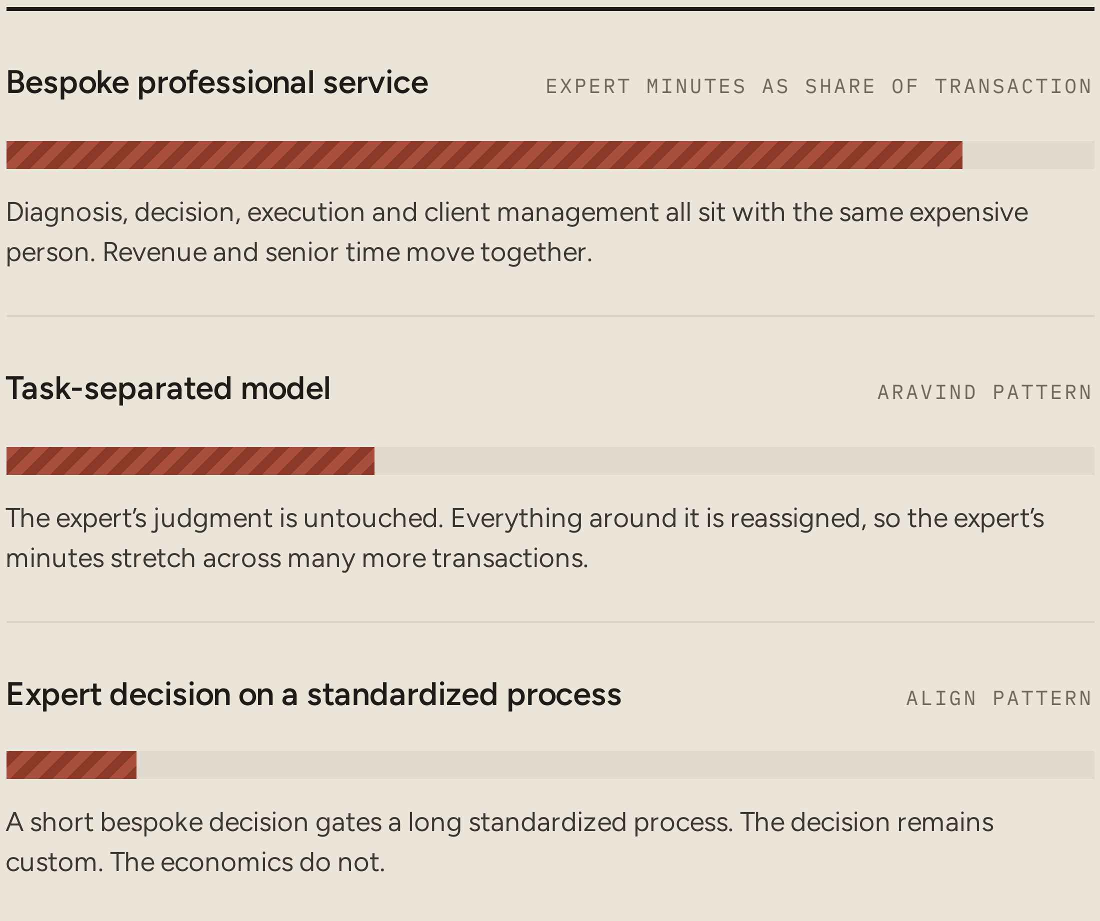

The proportions show the structural difference between these models. They are not drawn from data.

This reframes the question for consulting, agencies, legal work, recruiting and implementation services. The useful question is not whether bespoke expert judgment is present. It usually is, and often it should be. The question is how many expensive expert minutes each transaction consumes relative to what the transaction is worth.

A fifteen-minute expert decision sitting on a high-value standardized process behaves very differently from ten hours of expert decision-making on the same transaction. Same expertise. Different business.

CASE

The standard lane

The most consistent single move in the historical record is not serving every customer better. It is being unusually precise about which customers the standard path will not serve.

Shouldice Hernia Centre takes primary inguinal hernias in comparatively healthy patients and declines recurrent hernias, large scrotal hernias and cases with significant comorbidities. An independent study of Ontario hernia patients over fourteen years found recurrence at general hospitals ranging from 4.79% at the highest-volume institutions to 5.21% at the lowest, against a standardized recurrence risk of 1.15% at Shouldice.

Some of that difference is repetition and technique. Some of it is selection, and the hospital does not pretend otherwise. Its own materials acknowledge which case types typically require approaches it does not use, and those are broadly the case types it does not accept.

MinuteClinic runs the same logic in retail primary care. Nurse practitioners and physician assistants work to written protocols built on national practice guidelines, across a deliberately short list of low-acuity conditions. Operators treat seasonal allergies and generally do not treat asthma. Chronic disease management sits outside the model. Referral relationships with local physicians are established before a clinic opens, because the excluded cases have to go somewhere.

TurboTax publishes the boundary as a number. Its free tier is scoped to simple returns, which the company describes as covering roughly 37% of taxpayers. Everything above that line was, for more than a decade, simply not served by the free product. The assisted human tiers arrived in 2017 and 2020, long after the software had scaled on the eligible subset.

Qmerit, which runs electric vehicle charger installations across a network of licensed electricians, publishes what a standard installation excludes: multi-story buildings without a dedicated meter, panel and service upgrades, load management devices. A load calculation happens before anyone quotes.

A NECESSARY CAVEAT

Exclusion does not remove complexity from the world. Shouldice's declined patients still need operations. MinuteClinic's chronic disease patients still need care. Exclusion keeps that complexity outside one company's standard operating system, and moves it to someone whose economics are worse. This is a firm-level solution, and it is worth being honest that it is not a system-level one.

What separates a real exclusion from a stated preference is whether it is written down and given to the buyer. Exclusions that live in a salesperson's head are not exclusions. They are things the salesperson will forget under quota pressure.

FINDING

Control beats prediction

Where uncertainty comes from matters less than whether you own the surface it arrives through.

A recurring difficulty in professional services is that the decisive variable often belongs to the client, and does not exist yet at the moment of sale. How responsive they will be. How clean the data turns out to be. Whether the new executive keeps the project. Whether anyone actually implements the recommendation.

No amount of diagnostic work resolves a variable that has not happened. But several businesses have found a way around it that does not involve prediction at all.

Karat runs third-party technical interviews. Candidate behavior is entirely uncontrolled. The environment is not: interviews happen inside Karat's own recorded video and coding platform, against a rubric calibrated once with the client, scored by employed interviewers rather than the hiring team. The company reports having conducted more than 600,000 interviews. Variance in candidates remains total. Variance in how that candidate is observed and scored is substantially constrained.

ServiceTitan does not control whether a plumbing contractor runs a competent business. It controls dispatch, estimating, invoicing and payments, which is the surface through which the contractor's operations pass. Frontdoor does not control appliance failure rates, but it does control claim intake and contractor dispatch.

Contrast the advisory position. A consultant can recommend that a client publish something, restructure something, hire someone. What happens next is on the other side of a wall the consultant cannot see through, and the consequences arrive as scope creep, delayed projects and blamed outcomes.

> Do you control the surface through which the customer's behavior reaches the work? Advising is not controlling. Implementing is.

STRATEGY

Sometimes the answer is not to productize the service

Several of the largest businesses built inside service categories never productized the service at all. They found a layer within it that was easier to turn into a product, and took that instead.

Trex is the clearest example, because nothing about deck building changed. Building a deck remains custom, local and judgment-heavy. Trex made the material excellent and built a certified contractor program around it, launched in 2001, which requires an application, an information session and three registered decks built to specification and reviewed by a company representative before a contractor qualifies. The program has grown from around seventy professionals to thousands across nine countries, and the company reports it has connected contractors and homeowners on more than 304,000 projects.

The homeowner still gets a custom deck with all its variance. Trex sells boards.

Semrush did not become an SEO agency. It productized the diagnostic and monitoring layer that agencies use, reporting 2025 revenue of $443.6 million on annual recurring revenue of $471.4 million across roughly 116,000 to 118,000 paying customers, with an acquisition by Adobe announced in November 2025. For context, industry survey data puts the average North American digital agency at around $4.43 million in revenue.

ServiceTitan did not become a plumbing company. It reported fiscal 2025 revenue of $771.9 million, up 26%, from roughly 9,500 contractor customers, with gross retention above 95%. Toast did not become a restaurant. It serves around 164,000 locations, and payments now generate the majority of its gross profit.

THE PROGRESSION THESE BUSINESSES FOLLOW

01 **Bespoke service**
The underlying work stays custom, local and expert-led. It does not change.

02 **Productized layer**
A tool, material or component inside the service becomes a standardized product sold to the practitioners.

03 **Captured workflow**
The tool absorbs adjacent steps. Scheduling, estimating, documentation, reporting. The practitioner now runs the business inside it.

04 **Captured transaction**
Payments and financing move onto the same rails. ServiceTitan reported a revenue mix at listing of 71% subscription, 25% usage-based financial services and 4% professional services.

05 **Proprietary data**
Running the workflow and the transaction produces a dataset nobody outside the platform holds.

06 **Automation of parts of the service**
The data makes it possible to automate steps the practitioner used to perform. The layer begins consuming the service it was built to support.

Worth noting what this does to the operators on the other side. A contractor running on one of these platforms gets better tools and, in the same motion, higher switching costs and a weaker bargaining position.

For an executive inside a service business, the question arrives before the productization question. Is there a layer inside this service that would be easier to turn into a product than the service itself, and is anyone building it yet?

CORRECTION

Narrowing buys margin. Modularity buys scale.

The obvious lesson from the exclusion cases is that productized services are narrower than the services they replace. Survey data on digital agencies points the same way, with firms that reduced their service offerings reporting markedly higher net margins than those that expanded. That is self-reported and an association rather than a demonstrated cause.

The counterexamples improve the finding. Datadog started narrow and broadened aggressively, and became more product-like doing it. The company reports that 83% of customers use two or more of its products, 49% use four or more and 26% use six or more. Adding a product did not require adding a delivery method, because the products sit on shared infrastructure. Atlassian shows the same pattern while explicitly prioritizing larger customers, with gross margin improving over the period rather than eroding.

So narrowing is not itself the mechanism. A stable configuration space is, and narrowing is the cheapest way to create one when you have no infrastructure. Modularity is the expensive way, and the only one that allows expansion later without a return to bespoke delivery. An agency that narrows to one service and one customer type may improve its margins considerably and stay small. That is a legitimate outcome, and it is not the same outcome as building something that scales.

THE PATTERN

Five ways to commit to the economics

Each of the cases above resolves the same problem. The provider needs to commit before knowing. There are five recurring ways to get there, and most companies consider only the first.

01

Observe

See the variable that determines cost before committing economically. A physical site assessment, a credit and income pull, a structured diagnostic, a load calculation.

HISTORICAL FORM

Qmerit runs a load calculation before quoting. **Desktop Underwriter** pulls the data that decides the answer. The failure mode is pricing a job you have not looked at, visible at **Sungevity**, which used virtual site assessments with no truck roll before its 2017 bankruptcy.

02

Control

Own the interface through which uncertainty reaches the work, so that unpredictable behavior arrives as structured, measurable input rather than as a surprise.

HISTORICAL FORM

Karat runs the interview inside its own environment. **ServiceTitan** owns dispatch, estimating and payments. **Frontdoor** owns claim intake and contractor dispatch. Advisory businesses generally own none of it.

03

Pool

Let individual outcomes vary while making cohort economics predictable. Requires scale, a short enough contract to reprice, and some constraint on adverse selection.

HISTORICAL FORM

Frontdoor at roughly $2.09 billion revenue and 55% gross margin on a cost base it cannot forecast per member. **Zillow** also held a portfolio and could not reprice the houses it had already bought.

04

Exclude

Define what fits the standard configuration, publish it, and route everything else to a different path or a different provider.

HISTORICAL FORM

Shouldice on hernia type and comorbidity. **MinuteClinic** on condition. **TurboTax** at roughly 37% of taxpayers for its free tier. **Qmerit** with a published exclusion list.

05

Price the outcome

When an attempt costs almost nothing and success is cheap to verify, charge for successes and absorb the failures. The provider stops needing to predict the job.

HISTORICAL FORM

Long established in contingent legal work and recruiting. Recently extended into software-delivered support, where **Intercom** charges per resolved conversation and reports more than 40 million resolutions.

Note what is absent. Automation is not a route. Nor is packaging, nor standardizing deliverables, nor adding software. Those are things companies do while pursuing one of the five, and they are frequently mistaken for the destination.

The failure cases are consistent about this. Atrium raised $75.5 million to run a law firm alongside a technology company and priced in subscriptions, on the assumption that the software would reduce marginal cost. It shut in March 2020, its founder describing the company as having failed to deliver better efficiency than a traditional firm. ScaleFactor raised over $100 million selling automated back-office accounting and closed the same year, after reporting established that the work was being done manually by accountants in Austin and an outsourced office in the Philippines. Bench, a subscription bookkeeping business with more than 12,000 customers, shut abruptly in December 2024. Each fixed the price before fixing the cost.

THE CURRENT QUESTION

What AI changes, and what it does not

The reasonable hypothesis is that if understanding a customer's situation becomes cheap enough, customization itself becomes scalable, and the historical constraint dissolves. The evidence supports a narrower conclusion.

It makes observation cheaper

AI can inspect documents, websites, transaction histories, alert streams, codebases and support conversations at a cost that would have been prohibitive a few years ago. Diagnostic labor falls, and the share of customers who can be brought into a standard path rises with it.

The constraint moves rather than disappearing. Published resolution rates for AI support agents vary widely, and the variation tracks the quality of the customer's own documentation more than the capability of the model. One independent test across four small businesses found the two clients with comprehensive help documentation resolving roughly half of tickets, and the two with sparse documentation under a third. Cheaper observation raises the ceiling of the standard lane. It does not remove the lane, because the thing being observed still belongs to the customer.

It makes judgment cheaper per transaction

Upstart is the strongest documented case. The company reports that 91% of loans on its platform were fully automated in 2025 with no human intervention by Upstart, against roughly 70% at the time of its December 2020 listing, on a metric defined end to end from the initial rate request through to funding.

The point is not that underwriting has been solved. Lending partners retain control over their own credit policy and risk appetite, the models are designed and monitored by people, and roughly one loan in eleven still involves a human. The point is that per-transaction human judgment can become very small while policy, model design, monitoring and exceptions remain firmly human. Advertising shows the same movement from a different angle. Automated bidding has largely displaced manual bid management, and practitioners describe the new job as supplying the platform with better inputs rather than adjusting bids by hand. The judgment moved up the stack, and the platform captured it.

91%	40M+
Share of Upstart loans fully automated with no human intervention, 2025, against roughly 70% at its 2020 listing	Conversations resolved by Intercom's AI support agent, billed per successful resolution
COMPANY DISCLOSURE	COMPANY REPORTED

It makes continuous adaptation cheaper, where the provider owns the sensing surface

The most interesting possibility is that continuous re-observation could substitute for stability altogether. Rather than observe, classify and standardize, a system might observe, decide, execute, observe again and adapt indefinitely.

Large advertising platforms genuinely operate this way, and the reason they can is that the variables being re-observed live inside their own infrastructure. Nothing has to be reported to them by a client who may be slow or absent. Where that condition fails, the loop degrades to the quality of whatever the customer supplies, and the economics tend not to move. One large cybersecurity vendor has spent two years shipping agentic triage into a managed detection service while its disclosed subscription gross margin has changed by roughly a percentage point.

It extends where outcome pricing can work

Outcome pricing is not new. Contingent fees and placement fees long predate any of this. What AI changes is the cost of an unsuccessful attempt.

When inference is close to free, a provider can charge only for successes and absorb the failures without meaningful cost. That is what per-resolution pricing in customer support amounts to. A failed attempt no longer creates the same unit-economic problem, because the provider does not bill for it and running it cost almost nothing. The cost of a low resolution rate lands instead on the customer, in tickets they still have to staff.

The instrument has a clear boundary. It works where attempts are cheap, success is objectively verifiable and failure is low consequence. A failed support answer costs a redirect. A failed medical or legal judgment does not have that property, and this mechanism should not be assumed to travel there.

THE SUMMARY POSITION

AI has made observation, judgment and adaptation dramatically cheaper, and has widened the range of services where outcome pricing is viable. It has not produced a sixth route. A provider still has to be able to commit to the economics before knowing the job, and intelligence applied to delivery does not substitute for that.

APPLICATION

Ten questions for your own business

These are diagnostic rather than prescriptive. The answers tend to cluster, and the cluster tells you which of five states you are actually in.

01 **What single variable most determines what this customer will cost us?**

Name the specific fact. Not complexity, not fit. If you cannot name it, the rest of the diagnostic will not work.

02 **Can we observe it before we commit economically?**

And who has to do the observing. If it requires your most expensive person, the diagnostic is the service.

03 **Will it still be true in month three?**

A site does not change. A client's priorities, staffing and data quality do. Variables that are produced during delivery cannot be resolved by better discovery.

04 **If we cannot observe it, can we predict it reliably across a large cohort, and reprice before errors compound?**

Both halves are required. Pooling without repricing is a position, not a business model.

05 **Do we control the interface through which customer behavior enters the work?**

Recommending is not controlling. Running the workflow is.

06 **Which cases should explicitly not enter the standard path, and is that written down for buyers?**

Exclusions held informally get overridden at the end of the quarter.

07 **How many expensive expert minutes does each transaction consume, relative to what it is worth?**

This replaces the question of whether the work is standardized. Bespoke judgment is fine. Bespoke judgment occupying most of the transaction is not.

08 **Are we standardizing the economics, or only the deliverables?**

Four articles a month is a deliverable. Knowing why this client needs articles is the economics.

09 **Is there a layer inside this service that would be easier to productize than the service?**

Ask before attempting the harder version. Someone is usually already building it.

10 **Has AI reduced our judgment requirement, or only made execution cheaper?**

If the same person still decides what should happen for each customer, you have a faster service rather than a different one.

READING THE ANSWERS

Bespoke service	Packaged service	Productized service	Product-like layer	Should stay bespoke
The cost driver is neither observable nor poolable, and expert time dominates the transaction. Price it accordingly and stop apologizing for it.	Deliverables and prices are fixed, but the provider still cannot reliably commit to the underlying cost or expert effort. This is the state most companies mistake for the next one.	One or more of the five routes is genuinely in place, the standard lane is published, and expert minutes are a small share of the ticket.	The service remains bespoke and something inside it has become the product. Frequently the strongest position available.	High expert time, unpoolable cost, consequential failure. Some work is worth doing this way. The error is charging for it as though it were not.

CLOSING

The variability was never the point

For most of business history, turning a service into a product meant reducing variability until the work became repeatable. That is one route, and it is the one that gets attempted first because it is the one that feels like progress.

The historical record contains four others. Variability can be observed before you commit to it. It can be routed through an interface you own, so that it arrives as data rather than as a surprise. It can be pooled across enough customers that the individual case stops mattering. It can be excluded outright, which pushes it onto somebody else. And where an attempt costs nothing, it can simply be absorbed and billed only when it works.

None of these requires removing the expert. Several of the most product-like businesses in this study keep a highly paid specialist making a genuinely bespoke decision on every single transaction. What they have done is arrange matters so that the decision is short and the process around it is long.

And there is a further possibility, which the companies that took it rarely describe as a decision at all. The service can be left exactly as it is. Custom, local, expert-led, unscalable. What becomes the product is a layer inside it. The deck stays custom and the boards become a brand. The engagement stays bespoke and the diagnostic becomes a subscription. The trade stays a trade and the dispatch becomes a platform.

Before asking how to productize a service, it is worth asking whether the thing that wants to become a product is the service at all.

RESEARCH NOTES AND SOURCES

Figures are drawn from company filings and disclosures, regulatory records and contemporaneous reporting, with the source type noted below. Where a figure is company-reported and not independently verified, it is described that way in the text. Agency sector figures come from industry survey data and represent association rather than established cause. Independent testing of AI resolution rates diverges materially from vendor-reported rates, and both are noted above.

Frontdoor, Inc.
FY2025 annual report and quarterly results. Revenue, gross margin, adjusted EBITDA margin, renewal share, claims accrual as a critical estimate.
FINANCIAL DISCLOSURE

Zillow Group
2021 shareholder communications and contemporaneous reporting on the closure of Zillow Offers.
DISCLOSURE · JOURNALISM

Fannie Mae
Contemporaneous technical account of Desktop Underwriter, 1997; current Selling Guide provisions on lender underwriting judgment.
ACADEMIC · OPERATING GUIDANCE

Shouldice Hernia Centre
Independent fourteen-year study of Ontario hernia patients; peer-reviewed follow-up studies; hospital materials on case selection.
PEER-REVIEWED

CVS MinuteClinic
American College of Physicians position paper on retail clinics; peer-reviewed comparison of care quality for three conditions.
PEER-REVIEWED · SOCIETY POSITION

Align Technology
Annual reports on treatment planning and doctor control; October 2025 announcement on automated initial plan generation.
FILINGS · COMPANY RELEASE

Aravind Eye Care System
Published case narratives and clinical commentary on task allocation and surgical throughput.
ACADEMIC · JOURNALISM

Intuit and H&R Block
Published free-tier eligibility disclosures; product histories for assisted tiers.
COMPANY DISCLOSURE

Qmerit
Program documentation on standard installation scope, exclusions and network composition.
COMPANY DOCUMENTATION

Trex Company
Contractor program anniversary release, 2026, covering certification requirements and program scale.
COMPANY RELEASE

Semrush Holdings
FY2025 results and quarterly disclosures; announced acquisition, November 2025.
FINANCIAL DISCLOSURE

ServiceTitan and Toast
Listing documentation and reported results on revenue mix, retention and payments contribution.
FILINGS · SECONDARY COMPILATION

Karat
Company materials on interview volume, rubric calibration and delivery environment.
COMPANY DOCUMENTATION

Datadog and Atlassian
Quarterly results on multi-product adoption, customer cohorts, gross margin and net revenue retention.
FINANCIAL DISCLOSURE

Upstart Holdings
FY2025 annual report, including the definition and level of loan automation.
FINANCIAL DISCLOSURE

Intercom
Company statements on resolution volume and pricing; independent production testing of resolution rates.
COMPANY REPORTED · INDEPENDENT TEST

Atrium, ScaleFactor, Bench
Contemporaneous reporting on each closure, including founder statements and investigative coverage of automation claims.
JOURNALISM

Sungevity
Bankruptcy coverage, March and April 2017, including contractor accounts of remote site assessment.
JOURNALISM

Digital agency sector
Industry financial survey data on margin, service breadth and firm size. Self-reported and directional.
INDUSTRY SURVEY

A full source record with dated citations is available in the accompanying document.

PASTBEHAVIOR

PastBehavior publishes historical studies of how businesses have solved recurring structural problems. Claims are traced to primary sources where they exist, and marked where they do not.