

Webinar



SaaS and AI Pricing Report 2027 with Appinio

September 18th, 2026

Dr. Sebastian Voigt // sebastian.voigt@hy.co

Louise Leitsch // louise.leitsch@appinio.com

Nazeerah Balogun // nazeerah.balogun@hy.co

hy!



Dr. Sebastian Voigt
Partner & Co-CEO
hy Consulting Group

Director Axel Springer |
Director ProSieben |
Director Simon-Kucher |
PhD pricing | 100+ pricing
projects | “Pricing Friends”
podcast host



Nazeerah Balogun
Engagement Managerin
hy Consulting Group

>10 Pricing projects in the
field of software & IT |
Senior Consultant in Financial
Services & Pricing at d-fine |
Economic Consulting beim
DIW Econ | Inhouse
Consulting at innogy



Louise Leitsch
Chief Product & Research Officer
Appinio

10+ years as Behavioral
Scientists & Social
Psychologist | Now building
the sexiest research product
in the world | Expert in pricing
research (e.g. Conjoint, VWD) |
100+ conducted pricing
studies

Much of today's content is based on our analysis of **over 4,400 SaaS and AI companies**

SaaS & AI Pricing Report 2027

 consulting group ×  × 



Download our
**SaaS & AI Pricing
Report 2027** here!

<https://bit.ly/3SLCv2d>



We have more Pricing & Sales webinars upcoming in the next weeks

Register here for our upcoming
Pricing & Sales webinars!

<https://hy-consulting-group.webinargeek.com/channel/pricing-sales-webinars>



hy webinars



hy webinars × CARLSQUARE

Building a Better Exit Story Through Pricing & AI with Carlsquare

September 25, 2026

9am



Dr. Sebastian Voigt
Partner & Co-CEO, hy

Mark Miller
Founder & Managing Partner, Carlsquare

Eric Erbeck
Partner, Carlsquare

hy webinars × duwerk

Pricing for Humans & AI: Winning Customers and LLMs Alike with duwerk

October 02, 2026

9am



Dr. Sebastian Voigt
Partner & Co-CEO, hy

Katrin Kempe
CEO & Co-Founder, duwerk

Lili Frankus
CEO & Co-Founder, duwerk

hy webinars

From Stalling to Scaling: Rebuilding Account Management, CRM and Sales Routines for B2B Growth

November 06, 2026

9am



Frank Gehrig
Partner, hy

Anne Ringbeck
Senior Vice President, hy

hy webinars

The End of One-Size-Fits-All: 10 Value-based pricing Design Principles to Protect Margins

November 13, 2026

9am



Anne Ringbeck
Senior Vice President, hy

Dr. Amadeus Petzke
Partner, hy



Our agenda for today:

1. About hy + appinio
2. Key insights from the SaaS & AI Pricing Report 2027
3. Pricing transformation: How AI is reshaping key levers of monetization
4. Validation: How to validate the right AI pricing approach
5. Conclusion: Top learnings for your AI pricing strategy



2017

Founded

+70

FTEs

+600

Projects

+275

Clients

We support leaders in achieving their **ambitious growth targets.**

From Strategy to Execution: We Drive Results

We deliver actionable strategies — not slide decks that gather dust. Our senior team brings firsthand operational experience, ensuring that every recommendation is 100% implementable. We stay with you through the entire journey — from boardroom buy-in to on-the-ground execution — to unlock untapped value.

Entrepreneurs Helping Entrepreneurs

We understand what it takes to drive change in your organization, because entrepreneurship is in our DNA. Many of us are founders and investors, bringing this mindset to our clients. We know how to successfully tackle the unknown, driving the change and aligning with your ambitions and delivering transformative outcomes.

Unlock 10x ROI with Pricing Transformation

At hy, we don't aim for incremental improvements. We focus on transformation that drives measurable impact. By identifying and activating your biggest pricing levers, we target a 10x return on investment on every project within two years. For our clients, our support delivers short-term impact and helps to fund the journey.



Developing price improvement strategies for software and tech companies is our key DNA at hy





Driving value creation and pricing programs for private-equity owned companies is our home turf

Afinum

ALTOR

Armira[®]

AUCTUS

AURELIUS

BC PARTNERS

BBA
CAPITAL PARTNERS

BID EQUITY

Bregal
Unternehmerkapital

Bridgepoint

CINVEN

COMITIS CAPITAL

CVC

DEUTSCHE
PRIVATE
EQUITY

DRS
INVESTMENT

EGERIA

EQOS

Elvaston

EQT

FLEX
CAPITAL

EQUIVIA

GENUI

GREENPEAKPARTNERS

HANNOVER Finanz

H.I.G.
CAPITAL

IK Partners

KKR

LEA
PARTNERS

LIBERTA
PARTNERS

MAGUAR

MAXBURG
CAPITAL PARTNERS

Oakley Capital

palero

PARAGON
PARTNERS

PANGEA

PROM12

Rivean
Capital

Riverside.

spectrum
equity

SUMMIT PARTNERS

TYROL EQUITY

UFENAU
CAPITAL PARTNERS

WARBURG PINCUS



Rigorous, rapid research. Refreshingly human.

270+
Appinioneers

190+
Markets reached


25+
Nationalities

3,000+
World Industry-leaders
partnering with Appinio



365 SHERPAS	alrup	krivago	CAMPER	WE	home to go
St Hubert	granini	ABOUT YOU	Deloitte	nuna	FROSTA
share	BCG	Eurowings	* KICKBASE	LIVE NATION	Accenture Song
Coca-Cola	BASF	DB	playmobil	Jägermeister	sanotact

Appinio offers a flexible all-in-one solution




Powerful Technology

- Easy-to-use platform
- AI-enabled
- Collaborative features



Expert Guidance

- Advanced methodologies
- Dedicated research experts
- Flexible all-in-one experience



Global Panel

- Real world representation
- Data quality and rigor
- Real-time insights


Our agenda for today:

1. About hy + appinio
2. Key insights from the SaaS & AI Pricing Report 2027
3. Pricing transformation: How AI is reshaping key levers of monetization
4. Validation: How to validate the right AI pricing approach
5. Conclusion: Top learnings for your AI pricing strategy

We analyzed 4,400 OMR Reviews profiles, surveyed 153 SaaS companies, and spoke with 24 experts



Analysis of **4,400** OMR Reviews profiles to examine SaaS and AI pricing



153 responses from our online survey of SaaS companies regarding AI pricing strategies



24 interviews with external and hy-internal experts to generate in-depth insights into state-of-the-art AI pricing models

We spoke with many experts over the past few months: through projects, the Pricing Friends podcast, and dedicated 1:1 interviews for this report



Dr. Anna Lukasson-Herzig
nyris GmbH



Anne Ringbeck
hy Consulting Group



Arne Wolter
Appinio



Benedikt Böringer
Jamie



Christoph Röttgen
hy Consulting Group



Christoph Zöller
Instaffo



Fabian Heinrich
Mercanis



Florian Bunn
SCAYLE Commerce Engine



Gülsah Wilke
DN Capital



Ina Remmers
Rulemapping Group



Jenny Podewils
Leapsome



Julius Göllner
ARRtist



Malte Kosub
Parloa



Malte Landwehr
Peec AI



Mark Miller
Carlsquare



Marvin Müller
OMR Reviews



Melanie Riebel
Chaos



Nikolai Skatchkov
Circula



Robert Kaiser
Tacto



Stefan Peukert
Masterplan.com




Dr. Stephan Schmid
BU Bregal Unternehmerkapital



Dr. Toni Schmidt
FLEX Capital




Veronika Wax
Demodesk



Zoë Andreae
LECARE


Over the next 12 months, we expect six shifts in AI monetization



Product and Packaging

AI software will evolve from basic feature to add-ons into standalone products and product lines

With increasing autonomy, agents will no longer be merely integrated into existing packages but will be marketed as clearly defined services for specific tasks and workflows



Billing Model

Hybrid pricing is becoming the standard for AI in established software products

A fixed base fee ensures predictability, while variable components reflect actual AI usage and its costs



Pricing Metric

The seat is losing its role as the dominant pricing metric


By 2027, we expect a significantly broader use of tasks, workflows, and results as benchmarks for the variable price component



Agents and Outcomes

Agent pricing is establishing itself much faster than results-based pricing

Work can already be defined and priced today using tasks or workflows; performance-based compensation remains limited, for now, to use cases with clear measurability and attribution



AI Infrastructure

MCP is becoming the distribution standard; revenue comes from the services executed through it

Access via MCP often remains free; economic value is generated by actions, memory, runtime, and other infrastructure services that are actually used



Price Communication

Machine-readable pricing shifts from a marketing decision to a visibility requirement

Prices, pricing metrics, entry-level prices and package logic for enterprise SaaS are published in machine-readable formats, so AI systems draw on provider data rather than third-party sources


Our agenda for today:

1. About hy + appinio
2. Key insights from the SaaS & AI Pricing Report 2027
3. Pricing transformation: How AI is reshaping key levers of monetization
4. Validation: How to validate the right AI pricing approach
5. Conclusion: Top learnings for your AI pricing strategy


Today’s session uses an illustrative HR software company to explore the right AI monetization and validation approach

EXAMPLE OF AN HR SOFTWARE COMPANY TODAY

A platform managing core HR processes across the employee lifecycle, from hiring to payroll to performance:



TYPICAL MONETIZATION TODAY





What is happening: HR software providers are adding AI capabilities to every module at pace. New features range from workforce analytics to automated hiring workflows and AI-generated performance summaries



Why this is not straightforward: The same AI feature can serve as a free product enhancement, a premium add-on or a separate product entirely. The packaging decision shapes customer perception and long-term ARR from day one


↓

For each new AI feature, what is the right monetization approach and how do you validate that decision before committing to a model?

There are several options on how to include AI; today it is primarily treated as a feature and integrated into the core product

 INTEGRATION OF AI INTO THE PRODUCT PORTFOLIO

How are you currently integrating AI into your offering?
n = 153



 **Differentiated models are the exception:** add-ons, credits, standalone, and premium tiers combined account for just under 30%

 **One-fifth is not monetizing AI at all:** this applies almost exclusively to providers without AI at the core of their business model. Those who use AI centrally also have a monetization model



The market is split in two. One half integrates AI into the core product, while the other is spread across all other approaches

Source: SaaS & AI Pricing Report 2027 Survey



The diagram is a 2x2 matrix illustrating the relationship between the **Adoption of AI Features** (Y-axis) and the **Value of AI Features** (X-axis).

Y-axis: ADOPTION OF AI FEATURES
LOW (bottom) to HIGH (top)

X-axis: VALUE OF AI FEATURES
LOW (left) to HIGH (right)

Core Feature (High adoption - low value)

- AI job description generator**: Generate and optimize job postings in seconds
- AI writing & summaries**: Summarize meeting notes, draft documents etc.

Package Feature (High adoption - High value)

- AI workflow automation**: Automates onboarding, approvals & routine data tasks
- AI recruiting assistant**: Support recruiter across the hiring process

Experimental Feature (Low adoption - low value)

Add-on Feature (Low adoption - High value)


- AI sourcing & screening**: Source, screen and shortlist candidates automatically
- AI workforce planning**: Forecast workforce needs and simulate scenarios

Arrows indicate a progression from the Experimental Feature towards the Core, Package, and Add-on Features.

Package Feature: AI becomes the differentiator between base and premium package. The goal: Create a concrete incentive to upgrade

Does the AI solve a standalone customer problem and can it be used independently of the core product? If so, a standalone AI product may be the right choice

There are essentially four billing models, which differ primarily in the balance between predictability and cost risk



One-time payment at the time of purchase. Actual usage does not affect the price

Recurring, fixed fee per period (“flat rate”). No *direct* usage-based component


A recurring base fee is combined with a variable, usage-based payment

The price to be paid is entirely tied to the actual intensity of use


Payment		One-time payment	Recurring payment	Recurring payment	Recurring payment
Pricing Metrics		Usage-independent	Conditionally usage-based	Conditionally usage-based	Usage-based
Planned ability	Provider	Low (no ARR)	High (MRR/ARR)	Medium (base + variable)	Low (volatile revenue)
	Customer	High (one-time price)	High (fixed price)	Medium (fixed base, variable usage)	Low (highly variable costs)
Cost risk	Provider	Low (no ongoing costs)	High (for heavy users)	Medium (base cost predictable)	Low (margin always secure)
	Customer	Low (fixed price)	Low (fixed price)	Medium (risk of excess usage)	High (costs tied to usage)



Subscriptions continue to dominate the software landscape, but usage-based and hybrid billing models are gaining ground



A per-user price is not enough to reflect costs and willingness to pay; new metrics are coming into focus in software pricing



Industry Movement

HR Software →

Basic AI Feature

AI optimizes existing workflows. It increases product value and speed in the core product

💰 PRICING OPTION SPACE

✔ Included in the Core Tier

🔗 Higher-Tier Upsell

👤+ Add-on Flat Rate

AI Product

Standalone value-added feature & new module. Focus on advanced capabilities

💰 PRICING OPTION SPACE

🕒 Credit- or usage-based

🧩 Modular Add-on

💰 Hybrid (Base + Tokens)

Agent

Completes predefined tasks completely autonomously. Replaces individual operational steps

💰 PRICING OPTION SPACE

🔧 Per Job / Per Task

👤 Digital Worker Seat

⚡ Pay-as-you-go (PAYG)

Workflow Suite

Orchestrates complete end-to-end business processes across multiple agents

💰 PRICING OPTION SPACE

📈 Performance-based

📊 Volume Tiering

🤝 Hybrid Platform Fee

The pricing option space reflects the most common and conceivable models in the market



AI features improve quality and efficiency, but they are based on a “human-in-the-loop” concept; there are two pricing metric approaches


 PER-SEAT METRIC			 USAGE-BASED METRICS		
+	High revenue expectations & ARR predictability	Fixed surcharge per license/month (e.g., +€20/seat)	+	Direct margin protection (COGS linkage)	Expenses rise in exact parallel with API/infrastructure costs
+	Familiar SaaS model	High acceptance among B2B buyers; no budget uncertainty for customers	+	Fairness & low barrier to entry	Occasional users pay very little; trying it out is affordable
+	No barriers to adoption	Employees use AI without restriction and without fear of additional costs (“tax effect”)	X	Budget shock and billing variance	Customers fear fluctuating or uncontrollable monthly costs
X	Margin risk with power users	High LLM token costs from excessive use squeeze the gross margin	X	Usage Restraint	Users hesitate to make prompt entries so as not to “waste” credits

Ideal Use Cases:
Daily co-pilots & standard assistants with predictable basic usage (e.g., email drafts, CRM summaries)

Ideal Use Cases:
Computationally intensive specialized tasks (e.g., video generation, data analysis, infrequent batch processes)

Credits are particularly well-suited as a bridge between technical cost dimensions and the visible unit, especially for versatile AI applications

CREDIT MODEL



Figma has successfully introduced a credit system that includes AI credits and offers flexible top-ups on a subscription or pay-as-you-go basis



Since introducing seats with AI Credit limits, Figma boosted its gross margin by 2.5 ppt to 85% and achieved a net dollar retention of 136%

Starter	Professional	Organization	Enterprise
	MonthlyAnnual	Billed annually	Billed annually
Free Free limited access to Figma products	Full seat €16/mo + 3,000 AI credits/mo	Full seat €55/mo + 3,500 AI credits/mo	Full seat €90/mo + 4,250 AI credits/mo
	Dev seat €12/mo + 500 AI credits/mo	Dev seat €25/mo + 500 AI credits/mo	Dev seat €35/mo + 500 AI credits/mo
	Collab seat €3/mo + 500 AI credits/mo	Collab seat €5/mo + 500 AI credits/mo	Collab seat €5/mo + 500 AI credits/mo
Select plan	Select plan	Select plan	Contact sales

Usage insights show that 80% of large customers use at least one AI Credit per week

Add-ons

Customize your plan with optional add-ons

AI credits

Purchase additional AI credits for your whole team to share—enable pay-as-you-go for flexibility and/or subscribe to save.

[Learn more →](#)



HOW IT WORKS



Included credits per seat: Each license includes a monthly allocation of AI Credits, which are pooled to accommodate varying usage levels within the team



Flexible top-ups: If you need more, you can easily purchase additional credits (pay-as-you-go or credit add-on)



IDEAL FOR




Cost coverage for volatile usage: Providers whose AI costs depend on consumption and who want to link these costs directly to usage



Seat incentive: Providers where the credit allocation per license costs less than buying additional credits, which encourages purchasing more seats

monday.com bills AI separately from seats; with the included credit allowance rising by plan tier and consumption driven by task complexity



HOW IT WORKS

- Seats and AI usage are billed as two separate units, seats cover platform access, credits cover AI consumption
- Every AI capability across the platform draws from one shared credit pool, no per-feature add-on pricing
- Credit use scales with task complexity, not headcount

IDEAL FOR

- Include a predefined AI token allowance within every package
- Allow customers to purchase additional AI capacity through standardized top-up package

Source: [monday.com Pricing](#); Monday.com is a cloud-based work management platform



Workable introduced a credit model for its Workable Agent and offers flexible top-ups



Standard

BUNDLE RECRUITING AND HR TO SAVE UP TO 20%

\$299/month

\$3,588/year

MONTHLY

ANNUALLY - SAVE 20%

1 - 20 employees

Recruiting

Candidate sourcing suite

Applicant tracking system

Unlimited active jobs

5,000 AI credits

\$0.12/credit

\$600 total

Assessments

Workable Agent

Premier

BUNDLE RECRUITING AND HR TO SAVE UP TO 20%

\$599/month

\$7,188/year

PAY ANNUALLY

1 - 20 employees

Recruiting

Candidate sourcing suite

Applicant tracking system

Unlimited active jobs

10,000 AI credits

\$0.10/credit

\$1,000 total

Assessments

Referrals

Workable Agent

Enterprise

BUNDLE RECRUITING AND HR TO SAVE UP TO 20%

\$719/month

\$8,628/year

PAY ANNUALLY

1 - 20 employees

Recruiting

Candidate sourcing suite

Applicant tracking system

Unlimited active jobs

50,000 AI credits

\$0.095/credit

\$4,750 total

Assessments

Referrals

Workable Agent

Need more than 50,000 AI credits?

For larger bundles, our sales team can build a custom AI credit package with volume pricing tailored to your hiring.

Contact us for a quote



HOW IT WORKS



Included credits per account: Every plan starts with 3,000 pooled AI credits. Screening costs 1 credit, sourcing 2 and candidate chat 10 credits per interaction



Flexible top-ups: Bundles range from \$600 for 5,000 credits to \$4,750 for 50,000 credits, with automatic volume discounts



IDEAL FOR



Cost coverage for volatile usage: Teams only pay for actions the agent actually completes. No flat fee for unused capacity



Volume incentive: Credit prices fall from \$0.12 to \$0.095 per credit as bundle size grows, giving high-volume teams a clear reason to prepay

Source: [Workable](#)

27 | hy Webinar | SaaS & AI Pricing

Credit system design determines whether it captures margin or drives adoption, it depends

 CREDIT SYSTEM - DESIGN LEVER

	MARGIN-ORIENTED	BALANCE	ADOPTION-ORIENTED
ALLOWANCE LEVEL How much usage is included	Minimal - usage is paid from the start	Covers normal usage, top-ups for peaks	Covers peak usage, top-ups rarely needed
WEIGHTING What consumes credits	Every AI call consumes credits	Weighted by cost and complexity, simple actions free	Only executed work consumes credits
VALIDITY Unused credits	Expire at the end of each period	Included credits expire, purchased credits carry over	Remain valid
CONVERSION Action to credit	Not published, adjusted by vendor	Published, adjusted with notice	Published and contractually fixed
LIMITS Behaviour at zero balance	Usage continues, billed as overage	Paid actions pause, alerts before the limit	Hard stop, no additional cost

Margin capture (dynamic, vendor upside)

Adoption (predictable, fair)

The balanced approach protects margin without discouraging usage - AI stays included in the offering, while exceptional consumption is monetized rather than normal usage

It depends on the use case: agent services are priced completely differently depending on how it can be delineated in its offering

 **DIGITAL EMPLOYEE**



Unit: Prospects contacted per month, specifically per lead and not per activity

Price: Growth \$3,750 per month with an annual subscription, \$36,000 per year

Limit: 2,000 new prospects per month and 5 end users; price increase capped at renewal

Growth
Starting at \$3,750 / mo
Billed Annually
For teams putting a digital worker on their outbound pipeline for the first time.

- ✓ Up to 5 End Users
- ✓ 2,000 New Prospects per Month
- ✓ 11x Managed Gmail Mailboxes
Included: Domain Setup, Warmup, Inbox Rotation, Monitoring
- ✓ Bi-directional CRM Sync
- ✓ Onboarding + Support

Talk to sales →

Pro
Custom
Billed Annually
MOST POPULAR
For teams scaling outbound across markets, languages, and personas.

- ✓ Up to 10 End Users
- ✓ 5,000 New Prospects per Month
- ✓ 11x Managed Gmail Mailboxes
Included: Domain Setup, Warmup, Inbox Rotation, Monitoring
- ✓ Multi-language Outreach (105+)
- ✓ Bi-directional CRM Sync
- ✓ Tailored Onboarding + Support

Talk to sales →

Enterprise
Custom
Annual Or Multi-Year
For teams orchestrating outbound at scale across every segment and channel.

- ✓ Unlimited End Users
- ✓ 10,000+ Prospects per Month
- ✓ 11x Managed Gmail Mailboxes
Included: Domain Setup, Warmup, Inbox Rotation, Monitoring
- ✓ Bi-directional CRM Sync
- ✓ Custom Integrations + API
- ✓ Security and Compliance
SSO, Custom DPA, Custom SLA
- ✓ Dedicated CSM + FDE

Talk to sales →

 **USAGE/CREDITS**



Unit: Credits, tiered by task complexity

Price: Plus \$29.99 per user per month with 3,000 credits

Limit: Everyday Asks: 2 to 250 credits, Deep Work: 250 to 1,000, Big Builds: 1,000 to 2,500; the agent pauses if the limit is exceeded

Plus
Enough for everyday usage.
\$29.99 /mo per user
3k credits / user / mo
Standard usage
Try for free

Pro
More usage, more power.
\$99.99 /mo per user
15k credits / user / mo
5x more than Plus
Try for free

Max
For the heaviest workloads.
\$199.99 /mo per user
35k credits / user / mo
~12x more than Plus
Try for free

Enterprise
Everything in Max, plus the controls your security review will ask about.

- Everything in Max
- Shared usage + bonus credits
- Audit logs
- HIPAA compliance & signed BAA
- Dedicated support
- Onboarding & enablement

Talk to sales

Sources: [11x](#), [Lindy](#)



AI expands traditional pricing model development to include new questions about the product, usage, value, and costs

Step		Explanation
I	Objectives and Strategic Guidelines	A clear vision defines what the product and pricing model are intended to achieve. This also includes the strategic role of AI: Should it strengthen adoption and differentiation or be monetized as a standalone service?
II	Competitive Positioning	The competitive analysis shows how comparable offerings are positioned and monetized, and which providers are part of the consideration set from an AI perspective. Product scope, packaging, and pricing reveal relevant opportunities for differentiation
III	Product and Packaging	Product and price are inextricably linked. In addition to traditional features, it is necessary to define how AI features, add-ons, agents, or workflows will be integrated into and packaged as part of the existing product offering
IV	Billing Model	The billing model defines how customers generally pay. In addition to traditional subscriptions and pay-per-use models, hybrid models, credits, and quotas make it possible to combine variable AI usage with the necessary predictability
V	Pricing Metric(s)	The pricing metric determines the unit by which the price scales. With AI, multiple metrics can become relevant simultaneously, such as seats for software access, as well as credits, tasks, workflows, or results for AI services
VI	Price Level	The pricing level translates customer value and willingness to pay into specific prices. With AI, additional variable costs, varying usage intensities, and the potentially higher value of autonomously delivered services must also be taken into account
VII	Price Differentiation	Price differentiation makes it possible to systematically capitalize on varying willingness to pay and usage patterns. Packages, volume discounts, and scope of services create relevant opportunities for differentiation between traditional software and AI offerings
VIII	Price Enforcement	Clear rules for price approvals, special terms, and discounts ensure the consistency of the pricing model. A structured pricing enforcement process ensures that the developed pricing logic is consistently implemented in B2B sales



Three things to ask yourself when deciding on your AI pricing and validating your assumptions

Packaging

Which AI features do customers simply expect as part of the core product, and which capabilities drive enough value to justify a separate premium add-on?

Billing Model

How important is price predictability and where are customers comfortable with usage-based charges?

Price Metric

Does our target segment prefer to pay for AI access per seat, per bundled credit allowance, or per completed agent action?


Our agenda for today:

1. About hy + appinio
2. Key insights from the SaaS & AI Pricing Report 2027
3. Pricing transformation: How AI is reshaping key levers of monetization
4. Validation: How to validate the right AI pricing approach
5. Conclusion: Top learnings for your AI pricing strategy

You cannot price AI features until you know which ones are worth building, so the research runs in two steps

\$

FROM FEATURE VALIDATION TO PRICE VALIDATION



Our illustrative HR platform has 15 candidates on the roadmap: five core modules, five AI features and five agents

 THE VALIDATION BACKLOG OF AN ILLUSTRATIVE HR SOFTWARE COMPANY

NON-AI CORE MODULES	AI FEATURES	AI AGENTS
NO AI, ALREADY SOLD TODAY	ASSISTIVE, HUMAN STAYS IN THE LOOP	AUTONOMOUS, ACT ON THEIR OWN
Employee records & org structure	Job description generator	Sourcing agent
Time tracking & absence	Writing & meeting summaries	Screening & scheduling agent
Payroll interface & compliance	Employee HR copilot (policy Q&A)	Onboarding agent
Applicant tracking (ATS)	CV parsing & candidate ranking	Workforce planning & simulation
Performance reviews & goals	Performance review first drafts	Compliance monitoring agent

Every one of these 15 could end up as core, add-on, tier differentiator or a separate product.

Let's start this journey

Step 1: What should we build?

Kano, MaxDiff & TURF, applied to the same HR platform

Kano, MaxDiff and TURF, each answers a different question and you need all three answers

VALIDATING THE PRODUCT AND THE FEATURE SET		
KANO	TURF	Based on a MAXDIFF
What kind of feature is this? Two questions per feature: How you feel if it is there, and how you feel if it is not. The gap classifies every feature as must-be, performance, attractive, indifferent or reverse WHAT YOU DECIDE WITH IT What you are simply expected to have, and what you can legitimately charge extra for	Which combination reaches the most? Takes the individual-level MaxDiff results and searches for the set of features covering the largest share of customers, rather than the set with the highest average scores WHAT YOU DECIDE WITH IT Which three or four features together reach the widest base, which is the core question in bundling	What matters most? Respondents repeatedly pick the most and least important item from small sets. Forced trade-offs put every feature on one common scale, so nothing scores “important” by default WHAT YOU DECIDE WITH IT The priority order for the roadmap, and which features are strong enough to carry a package
ALWAYS FLANKED BY U&A QUESTIONS Who does this work today and at what cost · which budget line would pay for it · who decides and who can block · what the internal AI policy allows · which billing logic is understandable internally		AND ALWAYS SPLIT Run all three by ICP, persona and market. Averaging a mid-market recruiter and an enterprise works council produces a feature set that fits neither.

Kano tells you which AI features customers simply expect, which ones you can charge for, and which ones do damage


KANO CLASSIFICATION OF THE 15 ROADMAP CANDIDATES				
MUST-BE	PERFORMANCE	DELIGHTER	INDIFFERENT	REVERSE
Expected. Absence hurts, presence adds nothing	Linear. The more and the better, the happier	Attractive. Drives excitement and willingness to pay	No measurable effect either way	Satisfaction falls when the feature is present
- CV parsing & ranking - Writing & meeting summaries	- Employee HR copilot - Job description generator - Screening & scheduling agent	- Sourcing agent - Workforce planning & simulation - Compliance monitoring agent	- Engagement sentiment scoring - AI-written org-chart narratives	- Autonomous performance ratings - Attrition-risk flags to managers
PRICING IMPLICATION	PRICING IMPLICATION	PRICING IMPLICATION	PRICING IMPLICATION	PRICING IMPLICATION
Put it in core. Charging separately reads as a price increase, not as a new product	Tier differentiator or usage metric. Satisfaction scales with volume, so the price can too	Paid add-on or premium tier. This is where step 2 should spend its budget	Do not build it, and certainly do not pay inference costs for it	Remove, or make it opt-in with a full audit trail and human sign-off
Kano categories migrate: Today's delighter is next year's must-be, and for AI features that cycle runs in 12 to 18 months, so never lock a delighter into a permanent add-on price you cannot revisit.				

Illustrative results for the example HR platform



Plotted on how much they delight and how much their absence hurts, the candidates fall into four clear groups

KANO SATISFACTION / DISSATISFACTION GRID



HOW TO READ IT

- Must-be**
expected, absence hurts
- Performance**
more is better, linearly
- Delighters**
delights, drives willingness to pay
- Indifferent**
no measurable effect


REVERSE SITS OFF THIS GRID
Autonomous performance ratings and attrition-risk flags score negative satisfaction. They plot below the axis and should come out of the roadmap.

THE PRICING QUESTION
The distance between the two upper clusters is what you monetise. Performance features scale with volume, so they belong on a usage metric. Attractive features carry a separate price of their own.

Illustrative results for the example HR platform



The features with the highest scores do not automatically make the best bundle, because reach beats ranking




Illustrative results for the example HR platform, n = 600



Every feature you add to the bundle buys less reach than the one before it, and the curve flattens after three

INCREMENTAL REACH PER ADDED FEATURE



Three agents reach 76% of the base. The fourth and fifth together add 11 points.

HOW MANY FEATURES ARE WORTH BUNDLING

- Three features do most of the work**

The first three take you to 76% of the base, and each one opens a group the other two never touch. That is the bundle
- The fourth and fifth are expensive**

Workforce planning adds 7 points and the sourcing agent 4. Two full agent builds for 11 points of incremental reach
- So they belong in the add-on tier**

Low incremental reach combined with high individual value is the classic profile of a paid add-on. Step 2 is where you price it

Reach is not the only criterion. Margin and build cost decide the final cut.

Illustrative results for the example HR platform, n = 600



The same feature list gives three different answers by segment, and that is where the pricing opportunity sits

TOP PRIORITIES AND CONSTRAINTS BY ICP		
SMB 50–250 employees	MID-MARKET 250–2,000 employees	ENTERPRISE 2,000+ employees
WHO BUYS HR generalist, no dedicated recruiting team	WHO BUYS Talent acquisition lead plus HR operations	WHO BUYS HR ops, IT, procurement and the works council
TOP 3 PRIORITIES	TOP 3 PRIORITIES	TOP 3 PRIORITIES
Employee HR copilot	Screening & scheduling agent	Compliance monitoring
Onboarding agent	Sourcing agent	Workforce planning
Compliance monitoring	CV parsing & ranking	Audit trail & controls
WHAT THEY REJECT Sourcing agent, since they hire 15 people a year rather than 150. And credit accounting they have to administer themselves	WHAT THEY REJECT Per-seat pricing for agents that four recruiters run on behalf of the whole company	WHAT THEY REJECT Autonomous decisions without sign-off, and any bill they cannot forecast for the budget round
PRICING CONSEQUENCE Everything included, flat per employee. Simplicity beats precision; a variable bill costs more in admin than it saves in fees.	PRICING CONSEQUENCE Base fee plus credits. Volume-driven, and happy with variable spend as long as it is attributable to hires made.	PRICING CONSEQUENCE Committed volume with a hard cap. Predictability and control are features here, and they are priced as such.

Illustrative results for the example HR platform



Step 1 told us what to build

Step 2: what should it cost?

Conjoint, Gabor-Granger and Van Westendorp, applied to the same HR platform

Which pricing method you need follows from how the offer itself is structured

💰 CHOOSING THE RIGHT PRICING METHOD		
CONJOINT CBC / ACBC	GABOR-GRANGER DIRECT DEMAND CURVE	VAN WESTENDORP PRICE SENSITIVITY METER
USE IT WHEN Several interdependent elements, a tier architecture and a real competitive set	USE IT WHEN One clearly delineated add-on or agent, in a category the buyer already understands	USE IT WHEN A genuinely new AI service where the market has no reference price to anchor on
HOW IT WORKS Respondents repeatedly choose between complete, competing offers	HOW IT WORKS Purchase intent is asked at a sequence of prices, moving up or down until the respondent drops out	HOW IT WORKS Four questions per respondent: too cheap, cheap, expensive, too expensive
WHAT YOU GET Part-worth utilities, willingness to pay per feature, and a market simulator	WHAT YOU GET A demand curve, a revenue curve and a revenue-optimal price point for that single item	WHAT YOU GET An acceptable range and an indicative optimal point, enough to set the ladder for a Gabor-Granger
~15 min interview, n ≥ 300 per segment. Save it for the decision that matters most	Fast and cheap, but blind to competitors. Reliable for the shape of demand, optimistic about the level	Measures price perception, not demand. A range-finder, never a pricing decision on its own
Flanked by U&A: what they pay for this today · who owns the budget · how much monthly variance is acceptable · a comprehension check on the billing unit itself		

Conjoint lets you test the package, the billing unit and the price level in a single exercise

CHOICE-BASED CONJOINT: ATTRIBUTES AND LEVELS

Package	Core only · Core + AI Assist · Core + AI Assist + Agents · Agents standalone
Agent scope	none · screening · screening + sourcing · screening + sourcing + compliance
Billing model	all-in per employee/month · base + credits · base + per completed hire · pure pay-per-use
Included allowance	none · 1,000 credits · 3,000 credits · 10,000 credits
Controls & audit	full autonomy · approval required for outbound actions · full audit log + sign-off
Price	€6 · €9 · €12 · €16 per employee per month
The trick for AI: put the billing unit itself in as an attribute, at economically equivalent expected spend. Then you are not only testing how much You are also testing in which unit customers want to pay.	

WHAT COMES OUT

1 ATTRIBUTE IMPORTANCE

Price	31%
Agent scope	24%
Billing model	17%
Included allowance	14%
Controls & audit	9%
Support & SLA	5%

2 WILLINGNESS TO PAY PER AGENT

Screening & scheduling	+ €2.40 / employee / month
Sourcing	+ €1.70 / employee / month
Compliance monitoring	+ €2.10 / employee / month

3 MARKET SIMULATION


Share and revenue for any line-up, against the competitive set

Illustrative design and output for the example HR platform



A demand curve is enough for a single add-on, but with AI you also need the margin curve beside it

GABOR-GRANGER: SOURCING AGENT, PRICE PER CREDIT



HOW TO READ IT

Revenue-optimal is not adoption-optimal

Revenue peaks around €0.10 to €0.12, but adoption there is 68% against 41% at €0.14. If the agent is your land-and-expand motion, the revenue peak is the wrong target

AI has real variable cost

At €0.08 the pack is sold below inference cost for heavy users. Software pricing research traditionally ignores COGS because there were none. For agents the margin curve belongs on the chart

Van Westendorp comes first, not instead

With no reference price in the market it gave a plausible corridor of €0.09 to €0.18 and told us where to set the Gabor-Granger ladder. It never told us the price

Revenue peaks at €0.10–0.12, margin keeps climbing to €0.17. Classic SaaS research would stop at the revenue curve.

Gabor-Granger does not know your competitors exist, so validate the level against conjoint or real pilot conversion.

Illustrative results for the sourcing agent credit pack, example HR platform



Five things that break pricing research for AI, and none of them are about the price

 WHAT TO WATCH OUT FOR

1	Stated usage is not actual usage	Nobody can tell you how many credits they will need next quarter, because they have never bought credits before. Put the allowance into the conjoint as an attribute and validate the level against telemetry or a pilot
2	The buyer is not the user	The HR generalist loves the copilot, procurement signs the contract and the works council can stop the whole thing. Sample the decision unit, not just the enthusiast, and weight it the way the deal actually gets approved
3	Anything labelled AI attracts hypothetical enthusiasm	Stated interest in new AI features is systematically inflated. Competitive context inside the choice task and a switching gate afterwards bring the numbers back to something you can plan with
4	You may have measured comprehension, not preference	If a billing model scores badly, check first whether it was understood. A credit model that tests poorly is frequently a communication problem, and the fix is a different page rather than a different price
5	The answer has a shelf life	Kano categories move faster in AI than in any other feature class. Plan to re-run the short version every nine to twelve months, and write the contract so the price is allowed to follow

None of these are solved by a bigger sample. They are solved by a better study design, which makes the design itself a pricing decision.

Validating what to build and what to charge takes six to eight weeks from kickoff to a priced line-up

A REALISTIC VALIDATION PROGRAMME

WEEK 1	WEEK 2–3	WEEK 4	WEEK 5–6	WEEK 7–8
Hypotheses & long list Internal interviews, sales input, competitive teardown. Feature long list and ICP definition	Wave 1: what to build Kano + MaxDiff + TURF + U&A, n ≈ 600 across three ICPs, ≈ 15 min	Build & package decision What goes to core, add-on, tier differentiator or backlog, and what gets killed	Wave 2: what to charge CBC conjoint on the architecture, Gabor-Granger for standalone agents, n ≈ 600	Simulator & line-up Price level, metric, packages and the go-to-market story

Sequence	Validate the product before you price it. A conjoint on features nobody asked for returns a very precise price for the wrong product
Unit	With AI the metric is the real decision. Test seats, credits and outcomes against each other at equivalent expected spend, rather than only the price level
Segment	Run everything split by ICP. An averaged package fits nobody, and the distance between your segments is where the margin is

Our agenda for today:

1. About hy + appinio
2. Key insights from the SaaS & AI Pricing Report 2027
3. Pricing transformation: How AI is reshaping key levers of monetization
4. Validation: How to validate the right AI pricing approach
5. Conclusion: Top learnings for your AI pricing strategy

Structure your AI pricing to balance value with predictability, and validate the exact package and metric before you launch

HOW AI IS RESHAPING KEY LEVERS OF MONETIZATION

- 01 AI monetization starts with value delivered**
Not every AI feature should be monetized separately. Broadly adopted features can strengthen the core offering, while differentiated, high-value use cases create opportunities for upsell and add-ons
- 02 AI is pushing software pricing beyond the traditional per-seat model**
Seats remain relevant for human-centric software, but as AI takes over more work, usage, credits and actions become increasingly relevant metrics to capture value and costs.
- 03 Hybrid models can balance adoption, predictability and margin protection**
Combining a predictable base subscription with included usage and paid consumption beyond defined thresholds can encourage adoption while monetizing power users.

HOW TO VALIDATE THE RIGHT AI PRICING APPROACH

- 01 Validate what to build before you price it**
A conjoint on features nobody asked for returns a very precise price for the wrong product. Run Kano + MaxDiff + TURF first, pricing research second
- 02 The billing unit is the real decision, test it explicitly in the conjoint**
With AI, seats, credits, and outcomes all become viable metrics simultaneously. Include the billing unit as an attribute at equivalent expected spend
- 03 Run everything split by ICP, the segment gap is where the margin sits**
An averaged package fits nobody. SMB, mid-market, and enterprise differ in feature priorities, budget owners, and tolerance for variable billing

Read
our report.



<https://bit.ly/3SLCv2d>

Watch
our webinars.



bit.ly/hy-pricing-webinars

Listen
to our podcast.



hy — Your #1 partner for pricing and commercial excellence!