PRICING · A DOMAIN OF THE LEGAL OPERATING SYSTEM

Pricing & commercial models under AI pressure

Which work to reprice for AI, and how to defend the price of the rest

Derek Duffy
APAC Director, Alt-V Law · October 2026

© 2026 Derek Duffy. Licensed to Alt-V for authorised use.

This pillar sits inside the Legal Operating System

Strategy and Operations shape the firm; Pricing turns that shape into revenue

STRATEGY (APEX)	OPERATIONS	PARTNER PAY	PRICING	AI & TECHNOLOGY
Where to play and how to win. The partnership decision.	The economic engine. Converts strategy into margin.	The pay architecture. Pay cannot fix a shrinking pool alone.	How work is priced. Where AI pressure lands first.	What re-writes the hour count. Which forces this conversation.

A firm's pricing model determines which strategies it can execute, which clients it can hold, and which margin it keeps after AI touches the hour count. The pricing conversation is upstream of all three.

66% of firms are finding out about AI pricing through margin leakage

Headline numbers that frame every pricing conversation in 2026

34% / 66%

34% of firms have updated pricing for AI (BigHand 2025). The other 66% are discovering it through margin leakage.

84% / 23%

84% of firms offer AFAs. 23% of actual revenue is priced under one. The gap has not moved in 20 years.

90% HOURLY

90% of legal spend is still hourly. The pricing architecture has not caught up with the economics.

AM LAW $1-2k

Am Law 100 rates have crossed $1,000 for mid-tier partners and $2,000 for the top. The firm feels it in realisation.

ACC/Everlaw 2025: 52% of in-house teams using generative AI; 64% expect to reduce external reliance; 61% plan to push pricing changes to panel firms. The pressure is live and will not abate.

Five patterns that subtly cost firms their pricing position

Each is defensible on its own; the combination is why the margin story stays bad

RATE-CARD DEFENCE

The firm defends the current rate card because moving it feels risky. The erosion happens through realisation, not through the headline rate.

ONE-MODEL THINKING

The firm has a dominant commercial model and tries to apply it everywhere. Matter types that need a different model get the dominant one anyway.

GENERAL AI DISCOUNT

The firm concedes a general discount in response to a general AI-savings claim. No matter-type specificity; no shared-savings frame; margin given up permanently.

PANEL AS FEE SCHEDULE

Panel bids are responded to as rate-card submissions, not as commercial positions. The firm wins the panel at a margin it cannot sustain.

NO SCRIPT

No partner-wide answer to "why aren't you passing on the AI savings?" Three partners have had the conversation ad hoc, with no firm position behind them.

How small cuts in hours become a 13% revenue fall

The numbers do not need to be large to be material

Inputs

Modelled on the timesheets of ten large US firms. Toby Brown with Bruce MacEwen and Janet Stanton, Bloomberg Law, 16 May 2024.

- 5% fewer partner hours
- 20% fewer non-partner hours

Outputs

Revenue and margin both move more than the inputs suggest. The leverage runs the wrong way at the top of the P&L.

- 13% revenue hit
- About 7% margin hit
- Compounds across years as efficiency readings expand
- Invisible until the first panel renewal
- Impossible to recover without a pricing position

Not every matter type bends the same way

The right commercial model is different for each category

AI-HIGH

Document review, due diligence, discovery. 30–60% hour reduction expected 2–3 years. Move to fixed, capped, or outcome-based with AI efficiency priced in. Defending hourly here concedes margin.

AI-MEDIUM

Contract review, drafting, research. 10–25% reduction. The margin case is for hybrid structures - base plus performance or capped with efficiency sharing.

AI-LOW

Complex litigation, M&A, regulatory. <10% reduction. Hourly remains the honest default; the conversation centres on the value of the advice.

AI-HOSTILE

Trial advocacy, bet-the-company. Value-based or premium hourly. Position: this work was never a candidate for discount by efficiency.

Ten pricing models that depend on the work, and the role of AI

Each fits some matter types and breaks under others

Hourly	The default; the Toby Brown numbers quietly erode it. Honest for AI-low; progressively worse for AI-medium and AI-high.
Fixed	Transfers efficiency gains to the firm; requires scope discipline the firm may not have. Best for AI-medium and AI-high with scope controls.
Capped	Hybrid of hourly and fixed. The firm's hedge against both scope creep and realisation collapse.
Collared	Capped on both sides - floor and ceiling. Rare in practice; useful on volatile matter types.
AFA umbrella	Portfolio arrangement across matter types. Enables shared-savings framing with sophisticated clients.
Contingency / Success	Pure outcome. US-common; AU/UK context limited by regulatory frame and client appetite.
Retainer / Subscription	Recurring fee for a bounded service. Growing for in-house advice; the subscription-economics challenge is scope.
Value-based	Price keyed to client outcome, not firm effort. Requires a client dialogue many firms are Nascent on.
Portfolio / Panel	Panel-wide rates and volume commitments. Wins panels; compresses margin on the low end.
Hybrid (base + variable)	Alt-V's recommended target for AI-medium matter types. Base secures investment; variable rewards delivery.

Seven dimensions the Pricing Engine reads each matter type on

Output is a tiered price-band plus a recommended commercial model

Scope variability	How predictable the scope is. Low = fixed-fee candidate; high = hourly or capped.
Volume / repeatability	How often this matter type runs. High-volume supports productisation; low-volume stays bespoke.
Strategic value to client	What the client is actually buying. High-value supports value-based; low-value supports commodity pricing.
Risk profile	Technical and commercial risk the firm carries. Drives both the price band and the model.
AI impact category	AI-high / AI-medium / AI-low / AI-hostile. Determines which model category the matter sits in.
Competitive pressure	How much competition on this matter type. Intense pressure forces the fixed-fee or capped position.
Realisation	The firm's actual realised rate against the rate card. Low realisation drives the price-band recommendation downwards or into a different model.

Seven dimensions the Pricing Maturity Assessment reads a firm on

Nascent → Emerging → Established → Advanced → Optimised · the directional rule

Pricing capability and team	Is there a Pricing Lead? Pricing champions in practice groups? A dedicated analyst?
Commercial models used	How many of the ten models are in active use. Six-plus is typical mid-tier Established.
Data and analytics	Realisation visibility, matter-type profitability, lagged vs live. The data layer that drives everything else.
Governance and approval	Pricing committee, thresholds, audit trail. Weakness is usually enforcement at low-value matters.
Partner behaviour and discipline	Does the partnership actually live the governance frame. Scope creep absorbs 2–4% margin per quarter when it does not.
Client-facing pricing dialogue	The muscle to run the four-point script. Most frequently the gating dimension; nothing else lands without it.
AI-and-pricing readiness	Has the firm taken a formal position on sharing / retaining / hybriding the AI efficiency it has measured.

Three tools that work together

They compose - each reads a different layer of the pricing surface

PRICING ENGINE

Per-matter-type diagnostic. 7-dimension diagnostic → tiered price-band + recommended commercial model + four-point client script. Free tier shows the price band; full report runs through governance checklist and editable email template.

PANEL BID ECONOMICS

3-year P&L simulator for a panel renewal. Three AI substitution scenarios; margin bridge; sensitivity table; break-even analysis showing the panel discount at which year-1 margin hits zero.

PANEL BID BUILDER

7-dimension position-strength assessment against a specific panel brief. Strength / weakness map, gap register (pre-submission remediation list), win-themes keyed to the brief and the firm's challenger / incumbent position.

The Pricing Engine review for Thorne Marsden LLP

Invented firm · 180 FE · £35m fees · Category 2 commercial-advisory matter type

What the firm entered

Commercial advisory, mid-complexity. 7-dimension diagnostic.

- Scope variability L2 · Volume L3 · Strategic value L3
- Risk L3 · AI impact category 2 (medium) · Competition L4
- Realisation L2 (historical 83%)
- Current model: hourly · Target margin: 32%
- Firm size: standard (mid-market full-service)

What the Pricing Engine says

AI-medium, high-competition, mid-realisation. Hybrid (base + performance) recommended.

- Tiered price-band (low / target / high) generated against the firm's own cost model
- Recommended commercial model: Hybrid - the 70/30 base-and-variable structure
- Four-point client script generated with the firm's own figures
- Governance checklist + editable email template in the full report

The 3-year P&L assessment for the Q1 2027 panel renewal

Three AI substitution scenarios · margin bridge · break-even

What the firm entered

Standard rate card · mid-cap corporate GC panel · 3-year term

- Partner £950 / Senior associate £600 / Junior £400 / Paralegal £250
- Panel discount bid-in: 15% · Target firm margin: 32%
- Realisation 85% · WIP days 52 · Debtor days 58
- Aggressive AI scenario sensitivity: 1.4x
- Mix: commercial, employment, disputes

What the Economics says

Central scenario margin holds; aggressive scenario year-2 margin squeeze visible on hourly-structured mix lines.

- Central year-1: £1.23m revenue · £678k margin · 55.1%
- Aggressive year-1: £1.02m revenue · £559k margin · 54.8%
- Margin bridge: AI substitution absorbs the panel discount in Central; the Aggressive scenario sees the margin squeeze land in year 2
- Break-even at a panel discount of 63% against a bid-in of 15% - the firm has margin headroom

The position-strength verdict for the same panel

7-dimension assessment · bid readiness 45/100 · "At risk" if submitted today

TECHNICAL · IN LINE	PRICING · IN LINE	RESOURCING · BEHIND	TRACK RECORD · IN LINE	RELATION-SHIP · AHEAD	DEI/ESG · BEHIND	VALUE-ADD · BEHIND
L3 · Core sector depth adequate; competitors score similarly.	L3 · Fixed-fee readiness emerging; capped structures live.	L2 · Thin partner coverage on cross-border matter types.	L3 · Recent wins in commercial & disputes credible.	L4 · Existing relationships with 3 buyer-group entities.	L2 · Scored by this panel; firm's reporting lags.	L2 · Thin on secondments, academies, industry briefings.

Gap register surfaces three pre-submission remediation items (DEI reporting · value-add package · resourcing cross-border). Win-themes keyed to the challenger position: lead with Relationship, build case on Technical + Track Record, defuse DEI and VA in the response.

The defensible answer to “why aren’t you passing on the AI savings?”

The script reframes the conversation from discount demand to shared-savings design

1. Acknowledge the shift	AI has moved the economics on specific matter types; the firm has measured it and has a position. The concession is specific to those matter types.
2. Name which categories	Specificity defuses the general claim. AI-high categories carry a quantified efficiency; AI-low categories have not moved and will not.
3. Name what has been priced in	The firm has already adjusted on the matter types that have moved - through fixed fees, capped structures, efficiency-sharing clauses. Pre-empts the general discount demand.
4. Name the trade-off	Full pass-through transfers measurement cost, governance cost and pilot risk to the client. Shared-savings keeps the firm investing. A rate cut without that frame transfers all three.

Eight to twelve weeks, six phases, twelve templates

Lighter shapes compose from the same spine

Weeks 1–2 · SCOPING	T01 proposal · T02 data request pack · T03 interview guide (MP, CFO, PG heads, Pricing Lead, key clients)
Weeks 2–4 · DIAGNOSTIC	T04 Pricing Maturity Assessment (7 dimensions) · T05 matter-type portfolio worksheet (AI-suitability × competitive-pressure 2×2)
Weeks 4–7 · WORKSHOPS	T06 current-state (half-day) · T07 options-design (half-day, four commercial models shortlisted from the ten)
Weeks 6–9 · BUSINESS CASE	T08 structured business case · T09 benefits realisation model (margin uplift, realisation recovery, leakage reduction, AFA adoption)
Weeks 8–10 · BOARD	T10 board / partnership paper · partnership vote where required
Weeks 10–12+ · IMPLEMENTATION	T11 repricing cadence tracker (annual review + 7 trigger events) · T12 15–20 slide partner education deck

BOUTIQUE HEALTH CHECK

3–4 weeks against the maturity assessment alone. Scoping accelerator for sub-50 FE firms.

SINGLE PRACTICE GROUP

4–6 weeks on one practice group's matter types. The focused option when scope is clear.

COMPREHENSIVE REVIEW

12–16 weeks across all matter types with a per-practice-group workshop series.

What matter types at your firm are already being quietly discounted on an unverified assumption about AI?

And how would you know?

Derek Duffy · APAC Director, Alt-V · October 2026