

Pricing & commercial models under AI pressure

What I've built for firms deciding how to price work the AI has quietly re-costed, and defend the value of work it has not

General briefing · Alt-V Law

Derek Duffy, APAC Director, Alt-V

In short. 34% of firms have updated pricing for AI; the other 66% are discovering it through margin leakage. 90% of legal spend is still hourly; 84% of firms offer AFAs but only 23% of revenue is priced under one. The Toby Brown model: a 5% hour reduction on 20% of matter types produces a 13% revenue hit and a 7% margin hit. I've built a four-category AI-impact matter segmentation, a ten-model commercial taxonomy, three live tools (Pricing Engine · Panel Bid Economics · Panel Bid Builder), a seven-dimension pricing-maturity read, and an 8–12 week engagement with 12 named templates.

The frame

Not every matter type bends the same way.

The four-category AI-impact segmentation:

AI-high (document review, due diligence, discovery - 30–60% hour reduction, move to fixed / capped / outcome-based) · **AI-medium** (contract review, drafting, research - 10–25% reduction, hybrid structures) · **AI-low** (complex litigation, M&A, regulatory - <10% reduction, hourly holds) · **AI-hostile** (bet-the-company, trial advocacy - value-based or premium hourly).

Underneath the matter-type pricing sits the firm's capability to run pricing at all. Seven-dimension pricing-maturity read: pricing capability · commercial models used · data & analytics · governance · partner behaviour · client dialogue · AI-and-pricing readiness. Nascent → Optimised. The directional rule: a firm cannot run model-driven pricing from a Nascent data posture.

The four-point client script

When a client asks "why aren't you passing on the AI savings?":

1. **Acknowledge the shift.** AI has moved specific matter types; the firm has measured it.
2. **Name which categories.** AI-high moved; AI-low has not and will not.
3. **Name what has been priced in.** Fixed fees, capped structures, efficiency-sharing clauses already in place.
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.

The script reframes the conversation from "give us a discount" to "how do we share the gains while keeping the capability investment going."

The three tools

Pricing Engine. Per-matter-type diagnostic. Seven-dimension read; tiered price-band; recommended commercial model; four-point client script. **Panel Bid Economics.** 3-year P&L simulator on a panel renewal across three AI substitution scenarios; margin bridge; break-even. **Panel Bid Builder.** 7-dimension position-strength read against a specific panel brief; gap register; win-themes.

How the engagement runs

8–12 weeks, 6 phases, 12 templates. Scoping (weeks 1–2) · Diagnostic (2–4) · Workshops (4–7) · Business case (6–9) · Board (8–10) · Implementation (10–12+). Lighter shapes: 3–4 week Boutique Health Check; 4–6 week Single-practice-group review; 12–16 week Comprehensive Review.

Pairings and sample

Partner Pay: K-shape firms cannot fix pay without fixing pricing first (re-slicing a shrinking pool).

Operations / LOMM: matter-level economics is Pricing's data layer; a firm at LOMM Level 2 cannot run model-driven pricing. **M&A:** rate-card harmonisation after a merger is a direct Pricing Engine use case.

Thorne Marsden LLP: invented UK mid-tier (180 FE, £35m, 2025 AI pilot complete, Q1 2027 panel renewal approaching, firm is challenger). Sample pack carries all three tool outputs plus a seven-dimension Pricing Maturity Assessment.

What is the next matter type at your firm that is quietly being discounted on an unexamined assumption about AI - and how would you know?

Derek Duffy

APAC Director, Alt-V

NOTICE · FOCUS · CREATE