

DEMONSTRATION RUN: FICTIONAL COMPANY. Parsefield Ltd, its founders, customers, investors and every number below are invented. This document exists to show KAI's underwriting process end to end, in the exact format every real memo will follow. No investment decision was made or offered.

KAI · INVESTMENT MEMO #0

Parsefield Ltd · pre-seed · £400k round · decision memo

Date: Day 1 · **Verified as of:** Day 1 · **Engine:** KAI v0.1 (the first real memos run on v0.2) · **SLA:** applied → memo in 7h 18m · **Runtime:** frontier model under the FundNow specification

RUN LOG

Application to memo

STAGE	TIMESTAMP	ELAPSED
Application received (pitch and required information)	Day 1 · 09:02	0 min
KAI-SCREEN verdict: pass; evidence connections requested and received	Day 1 · 09:47	45 min
KAI-VERIFY table complete	Day 1 · 13:10	4h 08m
KAI-MEMO and decision issued	Day 1 · 16:20	7h 18m

SCREEN verdict (verbatim engine output):

```
{
  "verdict": "ADVANCE",
  "reasons": [
    "In scope: UK AI-native B2B software at pre-seed with verifiable revenue",
    "Identity, incorporation and product existence verified at T1/T2 on application",
    "Coherent 18-24-month follow-on story in a funded category"
  ],
  "kill_flags": [],
  "evidence_gaps": [
    "Counterparty confirmation of largest contract",
    "Reference calls pending"
  ]
}
```

1

Thesis

Parsefield sells agentic contract-intelligence to mid-market insurers: software agents that read, check and flag policy wordings and binders that today are reviewed manually by outsourced teams. Six paying insurers in nine months, verified £8.7k MRR growing about 2.1x per quarter, sold by a founding pair that unusually combines the technical depth (Imperial NLP

PhD, 8% university stake on clean USIT terms) with the market’s trust (eight years broking at Lloyd’s). The category has live budgets and visible follow-on appetite. Evidence quality is high for the stage; one material claim required correction (§3).

Composite	Decision	Team	Product	Market	Traction	Clock fit
7.15 / 10	BACK	8	7	7	6	7

Scores follow the v0.1 weights. Under v0.2, “Clock fit” is replaced by Capital & round dynamics, integrity becomes a gate rather than a weight, and the decision number is a calibrated probability against a published base rate.

2

Verification summary

Twenty-one material claims assessed: **12 VERIFIED (T1/T2) · 4 ATTESTED (T3) · 2 UNVERIFIED (T4) · 2 DISCREPANT · 1 PENDING**. Full table in the data room; the two discrepancies below, unabridged.

3

Discrepancy register

MATERIAL	D1. Revenue overstated 8.7%. Application claims £9,400 MRR . Stripe export (T1) reconciles to £8,650 : the £750 difference is a one-off pilot-setup fee from Halcyon Underwriting booked in the final week and presented as recurring. Founder response (logged): accepted the correction without dispute; stated the fee was “expected to recur” on pilot conversion. KAI treatment: memo uses £8,650 throughout; claim marked DISCREPANT→restated; no CRITICAL flag, because the underlying trajectory survives the correction. <i>This is what verification is for.</i>
NOTE	D2. Runway arithmetic. Application claims 7 months’ runway. Bank evidence (T2): £182k cash ÷ £31k monthly burn = 5.9 months . Founder counted a £40k Innovate UK grant instalment not yet received. KAI treatment: runway stated as 5.9 months verified, 7.2 months conditional on grant receipt (grant award letter sighted, T2; payment date unconfirmed → PENDING).

4

The company

Agents ingest policy wordings, binders and endorsements; check them against underwriting guidelines and regulation; flag conflicts with citations. Sold per-seat plus volume. Live product (used by KAI during diligence, T1), 140 active seats across 6 insurers, largest signed contract **£48k ACV** with Bridgewell Mutual (document verified T2; counterparty confirmation PENDING). Repo: 2,140 commits over 9 months, activity profile consistent with a four-person team building in production (T1).

5

Team

Dr. Amara Chen (CEO/CTO): PhD, NLP, Imperial College (thesis registry verified, T2); two first-author papers in the relevant sub-field; supervising professor's attestation (T3, quoted in data room): *"the strongest applied student I have supervised in a decade."* **James Okafor (CCO):** eight years placing specialty risk at Lloyd's (employment history cross-verified T2; two market references T3). Two engineers, full-time, verified. The rare pre-seed combination: the person who can build it and the person the buyers already trust. The unique knowledge: which specific clauses cost insurers money when missed, and what an underwriter will pay to never miss them.

6

Traction & economics

Verified MRR £8,650 (from £4,100 three months prior, about 2.1x a quarter, T1). Six logos, zero churn to date (9-month history, too short to lean on). ACV range £6k–£48k. Gross margin about 81% on current serving costs (T2, recomputed). Burn £31k/mo; verified runway 5.9 months → the raise is needed and honestly sized (£400k covers about 12 months at planned burn).

7

Market & the wedge

Mid-market insurers spend on manual wording review today; the displaced cost is real outsourcing invoices, not a hypothetical budget (two customer references confirm current spend, T3). Who dies first: the manual review vendors, not the incumbents' software. Wedge risk: horizontal document-AI platforms moving down-market; defence is workflow depth and the Lloyd's-trust channel, which generic platforms lack. Path to £10m revenue requires about 200 mid-market logos or Lloyd's syndicate expansion: plausible, not assumed.

8

The Clock plan

The 18–24-month question: can Parsefield reach a strong follow-on round? Verified trajectory (2.1x/quarter from a real base) reaching £40–60k MRR inside 15 months would place it firmly in range for UK/EU insurtech and applied-AI Series A appetite, which remains active in this category. Likeliest month-36 door: **secondary into an A/B round** (partial sale offered to backers, remainder retained). Exit door plausible but not underwritten (consolidators acquire in this space). Roll door available if trajectory is strong but round timing slips. Clock fit: 7/10.

9

Risks

1. **Concentration:** top logo (Bridgewell, £48k ACV, about £4k a month) = 46% of verified MRR; a Bridgewell churn resets the trajectory. Resolved by: counterparty confirmation + renewal evidence (requested).
2. **Category compression:** horizontal AI vendors bundling "good-enough" review. Resolved by: watching win/loss vs named platforms quarterly via telemetry.
3. **Founder-dependence in sales:** all six logos closed by Okafor personally. Resolved by: first non-founder sale within 12 months (tracked as a calibration prediction).

10

Calibration block: the public record

#	PREDICTION	P	RESOLVES
C1	Verified MRR $\geq$ £25k	0.45	Month 12
C2	Priced round at $\geq$ £12m pre-money	0.50	Month 18
C3	Company operating (not wound down/acqui-hired under cap)	0.80	Month 24
C4	$\geq$ 1 of current top-3 customers churned	0.30	Month 12
C5	First non-founder-closed sale	0.55	Month 12

In production every prediction is stated next to its base rate, fingerprinted in the public ledger on the decision day, then revealed and Brier-scored as it resolves; the company is named only with the founder's consent (v0.2 rule).

11

Decision & terms

DECISION: BACK, £120,000 offered to backers via SPV on the offered SAFE (£4.2m cap), joining alongside the named lead; each backer decides their own allocation. Decision conditional on: (a) counterparty confirmation of the Bridgewell contract; (b) restated revenue figure acknowledged in the round docs; (c) live Stripe telemetry connection maintained post-investment. Entry at £4.2m cap against verified £8.65k MRR is full but defensible given team quality and category heat; the discrepancy episode *raised* confidence in process integrity: the founder corrected without friction.

12

Decision record

DECISION	ISSUED BY	DATE	REVIEW REQUESTED
<i>(demonstration: no decision executed)</i>	KAI	n/a	n/a

Instrument note (v0.2 per-rail rule). This fictional deal is papered on a SAFE for demonstration. A real UK deal seeking SEIS/EIS would instead use an HMRC-compliant advance subscription agreement ($\leq$ 6-month longstop, no investor protections) or priced equity, because the YC SAFE's terms breach HMRC's advance-assurance conditions and would forfeit backers' reliefs. KAI v0.2 screens for this at intake.

This memo took KAI 7 hours 18 minutes, application to decision-ready. Your turn: **sign the NDA and drop your documents** for a decision within 10 days, on a **published SLA**. Backers: every memo like this, in full, on every deal.

What the engine does today

Autonomy is earned in public, never claimed. The state of KAI, step by step, updated as it changes.

STEP	TODAY	SHIPS NEXT
Claims: extracted from the application and matched to the evidence	KAI	—
Discrepancies graded; nine fraud checks run	KAI	—
Five dimensions scored; calibrated predictions; the full memo drafted	KAI	—
Evidence collected	HUMAN founders drop documents into the data room	Live read-only connections: payments, analytics, code, open banking
References and counterparty confirmations, in writing; Companies House checks	KAI automatic; no calls	—
The screen: through, or a no with reasons	KAI automatic; a person reviews any no on request	—
The final decision, with reasons	KAI automatic; a human can veto a yes, every override published; a person reviews any no on request	—
The calibration ledger	NEXT	Public: hash-committed predictions first, scores as they resolve
KAI Screen: the same verification, commissioned by an investor on a	NEXT	Once regulatory sign-off is in hand

company they are already looking at

Memo #0 ran on v0.1; the first real memos run on v0.2. The SLA has not yet been measured on a real application.

THE BASE RATES EVERY MEMO CITES

Seed round followed by a Series A within 24 months	15-20%
... before 2022	25-30%
Within 36 months	25-30%
Median time, seed to Series A	2.2 years
Bridge round at seed, Q1 2025	46%
Seed-funded companies shut down within five years	20-35%

Sources: Carta (graduation, timing, bridges, shutdowns); Dealroom (Europe, 36 months). Versioned with the engine specification; cited in full in every memo.

The record

THE RECORD

Scored against the base rate, not against a coin.

Every memo states one probability: a strong follow-on round within 18 to 24 months. Its fingerprint enters the public ledger on the decision day; the prediction is revealed and scored when the outcome is known. KAI recommends only at twice the base rate or better.

OPENS WITH THE FIRST DECISION

0

decisions made

0

predictions on record

—

skill v base rate

—

human override rate

Provisional until 30 predictions resolve, official at 100. Memo #0's five predictions are a demonstration, not scored.

KAI v0.1 · This demonstration memo follows the FundNow memo skeleton exactly. In production, sections 2–3 link to the full verification table in the vault, and the calibration block enters the public ledger as a dated fingerprint on the decision day, revealed and scored as each prediction resolves; the company is named only with the founder's consent. FundNow is a trading name of KAI Management Ltd, registered in England & Wales (company no. 17205149, registered office 71–75 Shelton Street, Covent Garden, London WC2H 9JQ), which is not authorised or regulated by the Financial Conduct Authority. Nothing on this page is an offer, solicitation, financial promotion or investment advice; the company described is fictional. [Privacy](#) · [Home](#)