Atlas

Open Questions & Next Experiments

Companion to the Atlas Investment Memo · Maya Okonkwo · 29 June 2026

How to use this doc. The memo is the static snapshot of the thesis. This is the working document that drives the next 30 days. Open questions are grouped by memo section; experiments are ordered by information-per-hour; the decisions log records why key calls were made so future-us can revisit them. The full budget build-up lives in the memo's The Ask section, not here.

1. Open questions

Grouped by memo section. Each has the question, why it matters, and the kind of evidence that would resolve it.

Problem & target customer

- **How large is real recoverable margin?** The entire pricing model rests on the 2-4 point hypothesis. Resolved by: median recoverable margin across the first 10-12 audits, measured, not modelled.
- **Is the daily user the finance analyst or the founder?** Changes the product surface and onboarding. Resolved by: 5 interviews with analysts (not just founders) about how they would actually use the worklist.

Business model & pricing

- **Will brands pay a share of recovery they have to verify?** Two operators flagged verification as friction. If it kills the performance fee, ACV drops materially. Resolved by: 3 paid pilots closed on the two-part price with no verification dispute.
- **What is the right performance-share window?** 12 months is a guess; too long feels punitive, too short undersells value. Resolved by: how the first pilots react to the term in contracting.

Go-to-market

- **Does the warm-network motion survive past customer 10?** Everything beyond 10 depends on a referral and audit-led motion we have not run. Resolved by: source of customers 11-20 and CAC on the first non-network deals.

Solution / wedge

- **Which leak categories are safe to auto-remediate?** Auto-apply is the path to the larger product but also the biggest trust risk. Resolved by: suggestion approval rate per category in suggest-only mode (target: a category above 90% approval before auto-apply).
- **How long does onboarding take across platforms?** If each brand needs bespoke integration, the model does not scale. Resolved by: time-to-first-recovery for brands #2-#5 on different stacks.

Right to win

- **Can we hire a founding AE who fits the audit-led motion?** Named as a top-three company risk. Resolved by: a signed AE with mid-market finance/ops sales experience by M5-M6.
- **Is the €70k/mo blended burn right once the AE is in?** The full build-up lives in the memo; what is still a guess is AE total comp (base + commission) and whether founder cash comp holds at below-market through M12. Resolved by: signed offers for the first two hires, which firm up roughly 80% of burn.

2. Next experiments (next 30 days)

Ordered by information-per-hour: highest-signal, lowest-cost first.

#	Experiment	Hypothesis tested	Success criterion	Time cost
1	Run 5 free 30-day audits from the warm list	Recoverable margin is 2-4 points	Median ≥ 1.5 points across the 5	~2 wks
2	Gate a paid tier on a signed recovery number with 3 of those brands	Brands pay a share of verified recovery	≥ 3 of 5 sign a paid pilot	~1 wk
3	Ship suggest-only remediation for the top 2 leak categories	Agent suggestions are trusted	Approval rate ≥ 90% in one category	~2 wks
4	Interview 5 finance analysts (daily users, not buyers)	The analyst will adopt the worklist	3 of 5 describe a fit to current workflow	~3 days
5	Onboard 2 brands on non-Shopify stacks	The wedge generalises across platforms	Time-to-first-recovery < 2 weeks each	~2 wks

3. Decisions log

Date	Decision	Rationale (revisit when…)
29 Jun 2026	Two-part price (platform fee + share of recovery), not flat SaaS	Aligns price with value and makes purchase near self-funding; revisit if verification proves contentious in pilots
29 Jun 2026	Start in suggest-only mode, not auto-apply	De-risks the agent and reveals which actions are safe to automate; revisit per-category once approval rate clears 90%
29 Jun 2026	SMB/mid-market wedge over enterprise	GTM velocity and warm-network access; revisit once retention and ACV data argue for moving upmarket
29 Jun 2026	Audit-as-lead-magnet as the core motion	The audit qualifies the buyer and produces the proof number; revisit if pilot-to-paid conversion stalls below 30%