Khive Labs pre-seed: our terms and the reasoning

$200,000 for 10% (implied $2.0M post-money, $1.8M pre-money), equity or SAFE. Open to discussion.

Memo for investors. Khive Labs, being incorporated as a spin-out of Khive (Khive Pvt Ltd, Rawalpindi / Islamabad, Pakistan). Written 23 November 2026. Facts about our work are current to 22 November 2026; market data carries its own date beside each figure.

This memo is not an offer to sell securities. Terms are discussed privately with qualified investors. Valuation methods used here are heuristics, not appraisals. We do not claim the company “is worth” any figure; we explain why we chose these terms.

1. Summary

- **What we ask:** $200,000 for 10% of Khive Labs, which implies a $2.0M post-money valuation.
- **What the market pays for this stage:** every public median we could verify for US pre-seed rounds is well above $2M post. The lowest is Carta’s median post-money SAFE cap of **$7.5M for rounds under $250,000** (Q2 2025), and caps have risen in every round-size tier since (Carta, Q3 2026). US seed rounds have a median pre-money of **$18.3M** in 2026, and **$21.0M** for AI companies (PitchBook-NVCA, data to 30 September 2026).
- **Why we ask less:** Khive Labs is pre-revenue, has one founder, mostly reports results on its own benchmarks, often on few seeds, and sits in Pakistan, where the whole startup market raised about $36.6M of equity in 2025 (Data Darbar). These risks are real and an investor should be paid for them.
- **Two angel heuristics, with inputs shown in the appendix:** the Berkus method gives $1.25M pre-money with Berkus’s own base caps and $3.75M with the caps he gives as an example for Silicon Valley; the Scorecard method on a US base gives about $6.4M pre-money. Our $1.8M pre-money sits at the low end of that range.
- **What $200,000 buys:** 12 months of work by two more ML engineers and the founder’s team, about $60,000 of compute (about 16 GPUs and 256 CPU threads around the clock for a year at the spot prices we pay), published results on outside benchmarks, a first game-studio pilot, and the incorporation.

2. Comparables: what pre-seed and seed rounds are priced at

All figures were read on 22-23 November 2026 from the sources linked. They are mostly US companies on each provider's platform. "Cap" means the valuation cap of a post-money SAFE, which is the usual pre-seed instrument; Carta reports that most SAFEs convert at their cap.

Source (date)	Figure
Carta, State of Pre-Seed Q2 2025 (21 Aug 2025)	Median post-money SAFE cap **$7.5M** for rounds under $250,000 (up from $6.5M the quarter before); **$10M** for $250,000-$500,000
Carta, State of Pre-Seed: 2025 in review (19 Feb 2026)	Median caps "hovered around **$10 million** for rounds in the $250,000 to $1 million range and **$15 million** for rounds in the $1 million to $2.5 million range"; caps rose across deal sizes in 2025
Carta, valuation caps Q2 2026 (25 Aug 2026)	SAFEs over $2.5M: median cap **$35M**, up 40% in a year; 75th percentile near $60M
Carta, State of Pre-Seed Q2 2026 (13 Aug 2026)	AI companies "captured half of all pre-seed dollars so far this year"; 90th-percentile cap on SAFEs over $2.5M up to $100M
Carta, State of Pre-Seed Q3 2026 (12 Nov 2026)	"The median post-money SAFE cap went up year over year for **every round size tier**"; $2.5M+ deals: median cap **$40M** (from $30M a year earlier); only 29% of seed rounds in Q1-Q3 2026 converted SAFEs below their cap
PitchBook-NVCA Venture Monitor Q3 2026 (data to 30 Sep 2026)	US median seed pre-money **$18.3M** in 2026 (2025: $16.0M); AI seed pre-money **$21.0M** (2025: $17.5M); median seed deal $3.0M
Carta data, Peter Walker (10 Jul 2026)	Software seed rounds in the prior six months: median **$24.3M** valuation on $4.1M raised, 18% dilution
Data Darbar, via Profit / Pakistan Today (6 Jan 2026)	Pakistani startups raised about **$36.6M of equity in 2025** across about 10 disclosed rounds (2024: $22.5M; 2021: $347M)

Sources:

- https://carta.com/data/state-of-pre-seed-q2-2025/SMARTPANTS_PRESERVED_131SMARTPANTS_PRESERVED_132

- https://carta.com/data/state-of-pre-seed-2025-full-report/SMARTPANTS_PRESERVED_135SMARTPANTS_PRESERVED_136
- https://carta.com/data/safe-valuation-caps-q2-2026/SMARTPANTS_PRESERVED_139SMARTPANTS_PRESERVED_140
- https://carta.com/data/state-of-pre-seed-q2-2026/SMARTPANTS_PRESERVED_143SMARTPANTS_PRESERVED_144
- https://carta.com/data/state-of-pre-seed-q3-2026-full-report/SMARTPANTS_PRESERVED_147SMARTPANTS_PRESERVED_148
- https://nvca.org/wp-content/uploads/2026/10/Q3-2026-PitchBook-NVCA-Venture-Monitor.pdfSMARTPANTS_PRESERVED_151 (charts “Median VC pre-money valuation ($M) by series” and “Median AI VC pre-money valuation ($M) by series”)
- https://carta.com/data/linkedin-archive/SMARTPANTS_PRESERVED_155 (post of 10 July 2026, “VC Startup Fundraising Benchmarks From 1000 Rounds”)
- https://profit.pakistantoday.com.pk/2026/01/06/startup-funding-in-pakistan-rises-to-36-6-million-in-2025/SMARTPANTS_PRESERVED_159SMARTPANTS_PRESERVED_160

What this says. The closest US comparable by round size is a SAFE under $250,000. Its median cap was $7.5M in Q2 2025 and has risen since. Our $2.0M is about a quarter of that figure. We found no AI-only median for rounds this small, so we do not quote one; the AI premium shows at seed, where AI companies are priced about 15% above the overall median (PitchBook-NVCA: $21.0M against $18.3M).

What we could not find. No public source we found gives a median pre-seed valuation for Pakistan or for “emerging venture markets” in 2025-2026. MAGNiTT’s early-stage reports are paywalled and the public summaries give deal counts and sizes, not valuations. AngelList’s public State of Venture summaries did not give dollar medians we could verify. We therefore treat the US figures as an upper reference, not as the price of a Pakistani company. The Pakistani market’s size (about $36.6M of equity in a year) is the clearest public sign of how thin local demand is.

3. Assets: what an investor buys a share of

Khive Labs owns no revenue today. What it has is a working learning library, measured results, and a fast, cheap way of producing more of them. Every number below has its limits beside it; we publish results that go against us too.

The technology. libbrain4 (“brain v4”) is a general learning library in C and CUDA that keeps learning while it acts. Its benchmark mind is 419,844 bytes; the fight mind is about 640 KB; the image mind about 1.9 MB. Research began on 3 November 2026 and brain v4 on 13 November

2026.

Results (from our fact sheet, each pre-registered: the pass or fail rule was written before the run):

- **Unseen kinds of task.** On 10 task families held out from training (our own suite), 0.51 ± 0.04 (spread between three independent trainings) against tuned PPO 0.076 and tuned RL² 0.068 (10 seeds each), after the same experience (77M steps). Caveats: our own benchmark; most of the gain is in a maze of forks (left/right choices) and in detours.
- **A fighting game against competition bots (FightingICE).** After 120 practice rounds against Iconic (the 2026 FightingICE LLM-track winner, a 1.5B-parameter language model) at full strength, it won the third 60-round life: 17 of the last 20 rounds, 32 of 60 overall; in the competition's best-of-3 format it won its final 8 games in a row. Caveats: one seed line; the game waits for both players (not the competition clock); the two practice lives were not won. Against ERHEA_PI (2020 champion), with memory kept between fights, it won 74-100 of 100 rounds in every fight from the second on, in 6 of 6 seeds (wait-for-both mode; ERHEA_PI is a 2020 search bot).
- **Reading unseen alphabets (Omniglot).** On the official evaluation alphabets, 86.6% at 5-way 1-shot and 96.4% at 5-shot, from 14x14 pixels with no pretrained features. Caveats: below published few-shot learners (prototypical networks about 98.8% on a related split); one eye.
- **Safety research.** In a study of guards that judge a self-improving learner's own rule changes, every guard design we tested let through changes that lost ability on task families the guard could not see (passers lose -0.027 life-sd, 3.7 se); a stricter floor guard made it worse on one family. Framing: selection pressure against measured scores in a small learner; nothing plans or deceives.
- **Speed of the code.** A timed raising of 32 minds went from 50.3 s to 3.6 s (14x) once its lives ran on all cores, with byte-identical results.

Live demo. One shared creature runs on our server at https://labs.khive.io/liveSMARTPANTS_PRESERVED_208, driven by brain v4 and taught by visitors in real time. A browser demo runs at https://labs.khive.io/demoSMARTPANTS_PRESERVED_210.SMARTPANTS_PRESERVED_211

The paper. A 17-page paper draft (20 November 2026, with a university co-author) describes what the library can and cannot yet do, including the failures. It is a draft, not yet submitted or peer reviewed.

Cost efficiency. All of this ran on one laptop GPU (RTX 3080 Laptop, 8 GB), our own shared server, and free cloud credits: our job ledgers estimate about $58 of Google Cloud trial credit and about $25 of AWS Activate credit used to 23 November 2026, about $83 in all (estimates;

disks and egress not counted). Single results cost cents: the larger Omniglot eye's training was about $0.31 of credit, the learned-gaze runs about $0.10.

Speed of iteration. Everything listed above was built and measured between 3 and 22 November 2026, by the founder working with AI research agents, while the parent company kept serving clients.

The parent. Khive (operating since 2017, incorporated with SECP on 24 June 2022) has about 50 engineers, roughly $30,000 a month of client revenue, and no outside funding. That revenue belongs to Khive, not Khive Labs; we do not count it in Khive Labs' value.

4. Use of funds and what $200,000 buys

Line (12 months)	Amount
Compute and cloud	about $60,000
Two ML engineers	about $100,000
First game-studio pilot and integration	about $25,000
Incorporation, legal, operations	about $15,000
Total	**$200,000**

What the compute line is in machine time. At the approximate spot prices on our roadmap (about $0.30 an hour for an L4-class GPU with its host, $0.007 per CPU thread-hour, $1.20 per A100-hour), $60,000 is about 16 GPUs and 256 CPU threads running around the clock for 12 months (about $42,000 + $15,700), plus about $2,400 for a 2,000 A100-hour benchmark block. It lets every claim be confirmed on more seeds and lets several lines run at once instead of in turn.

Milestones (targets, not promises; each judged by a rule written before its run, published win or lose):

1. **Published outside benchmarks, in the first month.** Our benchmark plan ranks them by our own estimate of the odds. XLand-MiniGrid against RL² re-run from its official code at the same steps (about $1-3 an arm; running now; our odds about 55%). Symbolic Alchemy against VMPO's published 155.4 (CPU, about $1-3; odds about 40%). bsuite memory and bandit axes against reference agents (under $1; odds about 70%). XLand at RL²'s full budget (about $10-15; odds about 15-20%). Atari 100k at 10 seeds (about 2,000 A100-hours, about $2,400); we expect to lose to BBF there and say so in advance.

2. **The all-rounder.** One mind raised on a mix of everything we test (unseen task families, fights, images, bandits, words) instead of task-specific minds. This takes most of the GPU line.
3. **The live creature, running all year,** and a public stream of it (one cloud GPU machine, about $250-350 a month at current spot prices).
4. **Hiring.** Two ML engineers, so the work no longer rests on one person.
5. **A first paying use.** One game-studio pilot for characters that learn, as the first test of the planned licence model (per product or per seat, plus paid integration support). We have no customer today.
6. **The company itself.** Incorporation of Khive Labs, transfer of the research IP from Khive, and investor paperwork.

5. Risks, stated plainly

- **Pre-revenue.** Khive Labs has no customers and no revenue. The licence model is a plan.
- **Small team.** One founder today. The research runs on the founder and AI research agents. Key-person risk is high until the two hires are in.
- **Mostly our own benchmarks.** The headline 0.51 is on our own task families. Outside benchmarks are only now being run, and we may lose on them.
- **Few seeds.** Several results stand on one to three seeds or one seed line (Iconic: one seed line; Omniglot: one eye). They may not hold up when repeated.
- **Results below the best published methods in places.** Omniglot is below published few-shot learners; Atari 100k is expected to be a loss.
- **Domicile.** Khive is a Pakistani company and Khive Labs' jurisdiction is not yet decided. A Pakistani company is harder for most foreign investors to hold, follow on in, and exit from; local startup funding is thin (about $36.6M of equity in 2025). Investors should take their own legal advice on cross-border holding.
- **Open publication.** Our Manifund agreement makes funded research and its code public, and we plan to publish code generally. The moat is the team's speed and the trained minds, not secret code.
- **IP not yet in the new company.** The work was done at Khive; it must be assigned to Khive Labs at incorporation.
- **Age.** The research is three weeks old. Fast progress so far does not guarantee it continues.
- **Competition.** Large, well-funded labs work on learning agents. We compete on size, cost and learning within one life, not on scale.

6. Why 10% for $200,000 is fair, and leans toward the investor

- **Ownership per dollar.** At the US median cap for rounds under $250,000 ($7.5M in Q2 2025, and higher since), $200,000 buys about 2.7%. At our terms it buys 10%, about 3.75 times as much.
- **The risks are priced in.** The gap between US comparables and our terms is our discount for the risks in section 5: pre-revenue, one founder, own benchmarks, few seeds, and the domicile.
- **It sits inside the angel heuristics.** Berkus gives $1.25M-$3.75M pre-money depending on the caps chosen; the Scorecard method gives about $6.4M on a US base. Our $1.8M pre-money is near the bottom of that range.
- **Low burn, measured output.** Results so far cost about $83 of free credits. The investor's money mostly buys people and machine time at known prices, not discovery of whether the code runs.
- **What it leaves for later rounds.** 10% at pre-seed leaves room for a seed round. If the outside benchmarks go well, a seed at even the 2026 US median ($18.3M pre-money) would mark this round up many times; if they go badly, we will publish that too.

We are open to discussion on the instrument (SAFE or equity), the cap, and the size.

Appendix: methods (heuristics, with stated inputs)

These are rules of thumb used by angel investors for pre-revenue companies. They depend on inputs we chose, shown here so a reader can change them.

A. Berkus method

Source: Dave Berkus, "The Berkus Method: Valuing an Early Stage Investment" (25 March 2012), https://berkus.com/the-berkus-method-valuing-an-early-stage-investment/SMARTPANTS_PRESERVED_388, and its update of 4 November 2016, https://berkus.com/25-years-updating-the-berkus-valuation-method/SMARTPANTS_PRESERVED_390. Five risk-reducing elements each add up to $500,000, for a pre-revenue value of up to $2M ($2.5M after rollout). The method applies to companies that could reach over $20M in revenue by year five. In 2016 Berkus wrote that the caps

should fit “the industry and the times”, gave $1.5M per element for a Silicon Valley “big data” startup and $500,000 for the same startup in Nebraska, and in a 2021 reply suggested dividing the local average pre-revenue valuation by five.

Element (risk reduced)	Our evidence	Share of cap	At $500k cap	At $1.5M cap
Sound idea (basic value)	a learner that adapts within one life, measured against tuned baselines; market unproven	70%	$350,000	$1,050,000
Prototype (technology)	working library, live creature, browser demo, pre-registered results	90%	$450,000	$1,350,000
Quality management team (execution)	one founder; parent with about 50 engineers; three weeks of fast output	60%	$300,000	$900,000
Strategic relationships (market)	AWS Activate credits, Manifund listing (no money yet), a university co-author; no customer or pilot	30%	$150,000	$450,000
Product rollout or sales (production)	none	0%	$0	$0
Pre-money total			**$1,250,000**	**$3,750,000**

Our $1.8M pre-money lies between the two. With no published Pakistani average, we cannot apply Berkus’s divide-by-five rule.

B. Scorecard (Payne) method

Source: Bill Payne, “Scorecard Valuation Methodology Rev 2019” (Angel Capital Association, 21 October 2019), https://angelcapitalassociation.org/blog/blog-scorecard-valuation-methodology-rev-2019-establishing-the-valuation-of-pre-revenue-start-up-companies/SMARTPANTS_PRESERVED_494. Take the median pre-money valuation of pre-revenue companies in the region and sector, rate the company against a typical one on weighted factors (100% = typical), and multiply.

Base (stated input): no regional median exists for Pakistan, so we use the US median cap for rounds under $250,000, $7.5M post-money (Carta Q2 2025), about $7.3M pre-money for a $200,000 round. This is a US base and overstates a Pakistani company’s price.

Factor (Payne's maximum weight)	Our rating vs a typical pre-seed company	Weighted
Strength of the team (30%)	75%: one founder, strong parent, no hires yet	0.225
Size of the opportunity (25%)	100%: game and device learning are real markets; unproven for us	0.250
Product / technology (15%)	125%: working, measured, cheap to run	0.188
Competitive environment (10%)	75%: large labs in adjacent work	0.075
Marketing, sales channels, partnerships (10%)	50%: no customer or pilot yet	0.050
Need for additional investment (5%)	125%: very low burn	0.063
Other (5%)	50%: domicile, IP transfer pending, open publication	0.025
Sum		**0.875**

0.875 x $7.3M = **about $6.4M pre-money** on a US base. Our $1.8M is about 28% of that. The difference is the discount for domicile and for what a US median does not capture.

C. What the three methods imply against our $2M post-money

Method	Pre-money	Post-money with $200k	Our terms vs it
Berkus, $500k caps	$1.25M	$1.45M	ours $0.55M higher
Berkus, $1.5M caps	$3.75M	$3.95M	ours about half
Scorecard, US base	about $6.4M	about $6.6M	ours about 30%
US median cap, rounds under $250k (Carta, Q2 2025)	about $7.3M	$7.5M	ours about 27%
Our terms	**$1.8M**	**$2.0M**	

Read plainly: only the most conservative Berkus reading prices us below our terms; every other reading prices us above them.