HIVEMIND

Pitch Deck (readable version)

Strategy is a luxury.

- Algorithmic and simulation-based strategic analytical frameworks have significantly improved over the past decade.
  - Simulation-Based Optimization, Strategic Agent-Based Modeling, Markov Decision Processes, Hybrid Optimization Models, etc.
- Thus far, using these improved frameworks in any valuable way has required expensive, expert R&D teams.
- These improved methods, however, have largely been inaccessible for small business owners and individuals, who don't have enough spare cash to spend on expert teams.

Strategy was a luxury.

- With the invention of generative AI, a door opened for the democratization of these frameworks.
- Most elite firms already use proprietary generative AI tools to automate the mapping of strategic frameworks onto their decisions. Generative AI tools can be used for this analysis.
  - But these firms face a net-cost in opening these tools up to the public: they'd lose their competitive edge, scaling is expensive, they have reputations to worry about, and they benefit from locking people into proprietary services instead.
- Possibility + reluctance = opportunity for a new player.

Hivemind's design

- Hivemind is a multi-stage algorithm that uses client data (e.g. input KPIs), online data (e.g. market trends), theoretical frameworks (e.g. Monte-Carlo simulations), and practical feasibility testing (e.g. regulatory compliance) to strategically analyze single decisions.
- It's designed to be adaptable to any decision, in any situation, for any type of client.
  - Hivemind can be used by a Wall-Street career enthusiast just as much as it can be used by a small restaurant in Chicago.
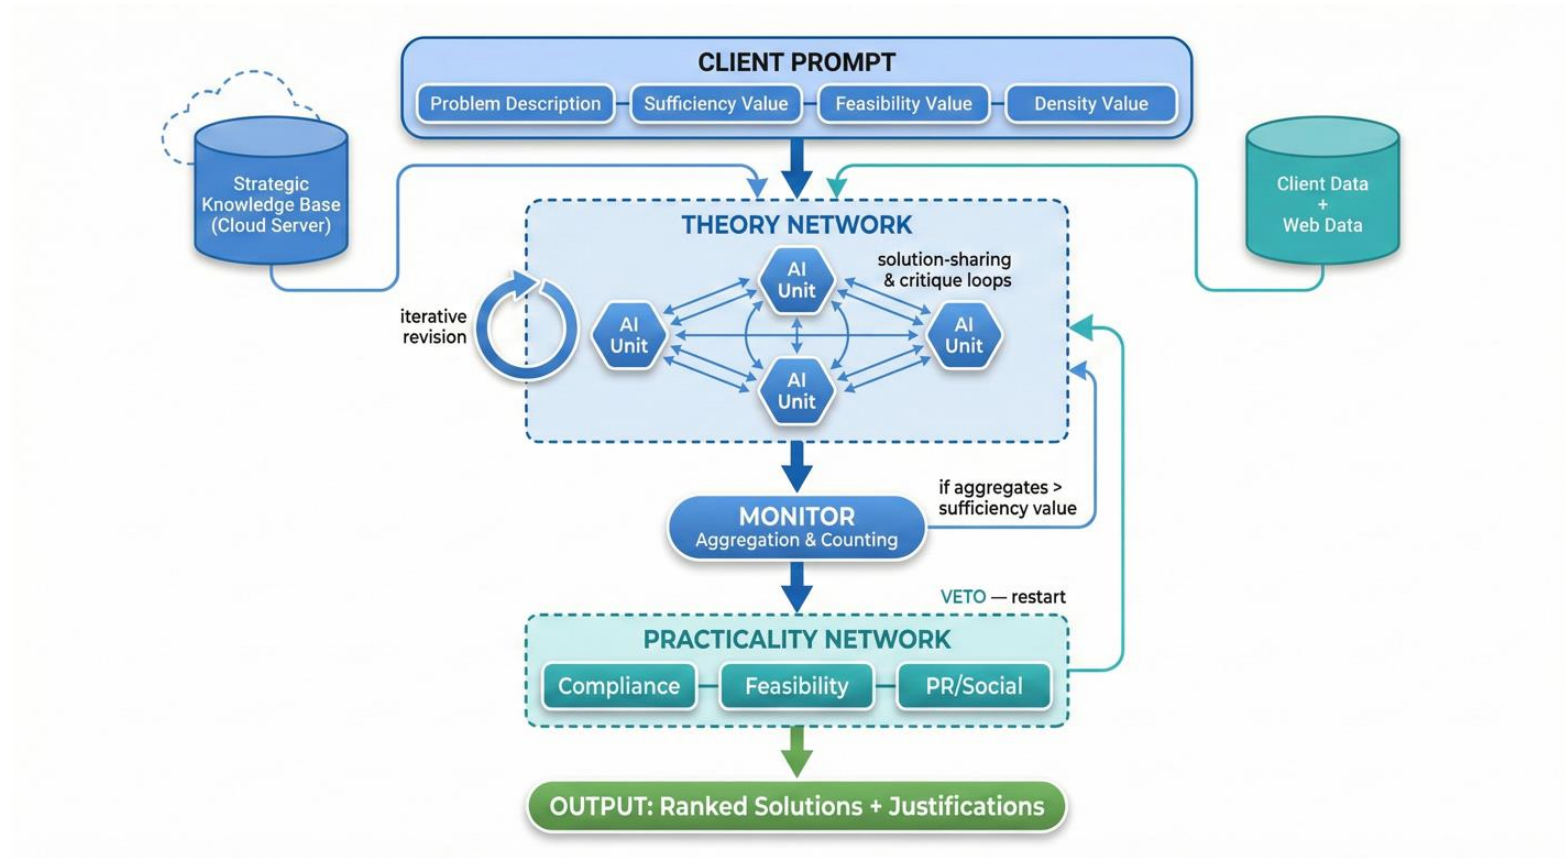
Hivemind's workflow (analogy)

- An analogy that explains hivemind's analytical process:

Hivemind gives you a team of AI strategists that debate your problem, then reality-checks their ideas before anything reaches you.

You describe your problem. AI experts, each drawing from different strategic theories, propose solutions, critique each other, and revise — back and forth until they converge on one or more similar strategic propositions. A referee groups similar ideas and decides when the debate is done. The survivors get scored by real-world advisors for feasibility. If scores are too low, everything gets thrown out and the experts start fresh.

Hivemind's workflow (visualization)

This is a visualization of Hivemind in its full complexity, not as the analogy provided.

Hivemind's value offer

- While generative AI (LLMs) are used to communicate the conclusions of strategic analysis to the client, the value lies instead in the frameworks used to conduct the analysis.
  - This is NOT a complicated ChatGPT wrapper. The researched and developed frameworks provide the value.
- The firm's structure will lean heavily into R&D to capitalize on this value differentiator. The idea is to have the firm function like a strategic think tank, where strategic theory is researched and developed natively.

Hivemind's ideal users

- Hivemind is designed for usage by small-scale agents (individuals, businesses) in highly competitive environments.
- Examples of ideal clients include restaurants, corporate employees (Wall Street types), startups, real estate agents, gyms, MBA students, private practices, etc.

Market opportunity

- TAM: There are 33.2 million small businesses in the U.S. alone. The global management consulting market is worth over $300 billion annually. Strategic advisory tools sit at the intersection of consulting and SaaS.

- SAM: Small businesses and competitive professionals who currently spend on consultants, coaches, MBA programs, or strategic advisory but are underserved by high-cost incumbents. Estimated at $15–25B in the U.S.

- First-year target: 500–1,000 paying customers in 2–3 niches.

Competitive landscape

Hivemind sits in a gap that no current player fills:

- High-end consulting has strategic rigor, but costs $500K+ per engagement. It's completely inaccessible to small businesses and individuals.

- Generative AI tools are cheap and accessible but have no structured methodology. Output quality depends entirely on the user's ability to prompt well. And forget about feasibility testing, multi-perspective debate, simulations, or framework depth.

- Business coaching and advisory is personalized but unscalable, of inconsistent quality, and rarely grounded in formal strategic theory.

Hivemind uniquely combines rigor with accessibility.

Go-to-market strategy

First small businesses, then individuals.

1. Start with small businesses in 2–3 verticals (restaurants, real estate, startups). Build industry-specific knowledge bases.

2. Use knowledge base learnings from initial verticals to enter adjacent markets. Each vertical enriches the platform for the next.

3. Once proven with businesses, adapt for personal decisions (career, negotiations, investments).

4. Direct outreach to small business associations and accelerators. Content marketing. Partnerships with business schools and incubators.

Business model

- One-time installation fee + recurring subscription. Tiered based on usage and analytical depth.

- Individuals (~$49/mo): career decisions, negotiation strategy, personal investment analysis. Limited queries per month, standard knowledge base.

- Small business tier (~$199/mo): competitive positioning, pricing, expansion planning, hiring strategy. Higher query volume, industry-specific knowledge bases, forward-deployed support.

- Enterprise/custom tier: dedicated forward-deployed engineers and academics, custom theory development, full data integration. Pricing on contract.

- Unit economics improve as the knowledge base scales; theory developed for one client enriches the platform for all clients.

Us

- Core/founding team: Michael Quintin, Ryan Pieper, Victoria Liu, Matthew Foster, Armaan Malhotra (all current UChicago undergrads)

Overall team structure:

- CS team: software engineers building infrastructure and forward-deployed engineers interacting directly with clients as personal software assistants.

- Research team: strategic theorists developing the knowledge base, plus forward-deployed academics interacting directly with clients to identify theory gaps and commission new research.

- Compliance team: regulatory experts keeping the practicality network current across industries.

Traction & milestones

Algorithm fully designed. MVP in progress, nearly completed. Talking to restaurant owners to research first product niche.

Next 1–3 months, shipping MVP to 10–20 initial customers in Chicago restaurant vertical. Validate value offer, collect feedback, and iterate on software's theoretical/practical analysis.

Next 4–8 months, expanding into another 2–3 verticals (researched needed here). Reach 200 paying customers. Iterate as above. Hire first academic and forward-deployed engineer. Begin developing the individual tier.

Next 9-12 months, launching the individual tier. Online and guerilla marketing, using Booth and other connections to reach competitive individuals.

Long-term competitive moats

- Proprietary LLM for privacy and accuracy: training an internal LLM on strategic reports will smooth the translation from client problems to analytical frameworks, as well as ensure complete data privacy.
  - At months 9–12, we may be able to start developing the LLM.
- Genuine theoretical progress: strategic theory developed at the firm, as well as the formalization of feasibility testing.

Both moats get stronger with scale. More clients means more data, better models, and richer theory.

The ask

$500K to build and launch Hivemind.

Funding allocation:

- Software = 40%: building basic software and infrastructure.
- Research = 25%: building the strategic theory and feasibility testing knowledge bases.
- GTM = 20%: first 200 customers. FDE hires. Marketing.
- Internal = 15%: firm-side compliance and admin.

Result = product-market fit, initial revenue, positioned to raise Series A.

Appendix 1: use case, small business

Maria owns a successful Thai restaurant in Austin and wants to open a second location. She describes the problem to Hivemind: "I have $200K saved, three neighborhoods I'm considering, and I need to decide where to open, when to open, and whether to keep the same menu or adapt it." The AI experts debate — some argue for the high-foot-traffic downtown spot using competitive positioning theory, others push for a cheaper suburban location using cost-advantage frameworks, one flags that her best-selling dishes skew toward a demographic concentrated in the east side. They critique each other, revise, and converge on two strategies. The real-world advisors then score both: the downtown option gets flagged for lease terms that could sink her if revenue dips in the first year. The suburban east-side location passes. Maria gets a recommendation with clear reasoning from multiple strategic angles, plus the knowledge that it survived a feasibility gauntlet — something that would have cost her $15,000 from a consulting firm.

Appendix 2: use case, individual

James is a VP of operations at a mid-size logistics company. He just received an offer from a competitor — better title, 20% raise, but a less stable company. He asks Hivemind: "Should I take this offer, negotiate with my current employer, or use it as leverage without intending to leave?" The AI experts weigh in from different angles — negotiation theory, career signaling, game theory on employer response, risk analysis on the new company's financials. They debate whether revealing the offer to his current boss strengthens or weakens his position long-term. They converge on two strategies: a structured counter-negotiation with his current employer using the offer as leverage, and a conditional acceptance of the new role with specific contractual protections. The real-world advisors flag that the leveraging play carries reputational risk in his tight-knit industry. James gets both options with honest tradeoffs laid out — the kind of advice you'd normally only get from a mentor who happens to have an MBA and knows your industry cold.