# THE NASH LAB PRODUCT PITCH:

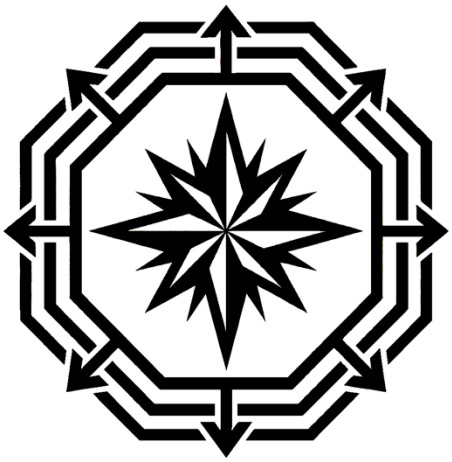

# HIVEMIND

Career success requires extremely strategic economic decision-making. Individual actors, if unable to make strategically optimal choices, suffer short and long-term consequences. Business schools attempt to fill this gap, educating economic agents on general principles conducive to career success, and rough frameworks typically accurate in their strategic conclusions. But even coming out of business schools, agents do not have any rigorous way of successfully applying the methods they learn, nor can learn all the strategic analytical frameworks at their disposal.

A truly successful agent would operate like a business expert — as if a strategic intellectual descended from the ivory tower to pursue a career in consulting or investment banking. There is, therefore, a knowledge gap between economic agents and business experts, and furthermore a need to democratize rigorous academia in the business strategy context. How can we give career-focused individuals and small organizations the ability to make strategic decisions like intelligent, educated, and rigorously refined business experts?

The solution is Hivemind, a strategic analytical tool combining theoretical and practical analysis, providing academic insight in pragmatic economic decision-making contexts.

Hivemind's workflow proceeds as follows:

1. Before the client prompts Hivemind, the tool will have access to any and all data the client has cleared for collection and usage by the model — as well as data scraped from the internet. This data can and likely will be used in its strategic analysis.
2. Before the client prompts Hivemind, a cloud server, updated and managed by admins inside the firm (using an admin software), will already contain a strategic knowledge base used later to analyze client problems. The strategic knowledge base will be a set of self-enclosed documents, each explaining a given strategic theory simply and clearly. In some of the knowledge base documents, certain mathematical or otherwise algorithmic simulations may be mentioned as necessary for analysis. These simulations will be stored separately from the knowledge base documents in a simulation base also hosted on the cloud server, written as python programs with clear inputs and outputs, accessible to be called on and input data into during later analysis.
3. The Hivemind workflow will contain four main components: the prompt, the theory network, the practicality network, and the monitor. These four components will work together (as follows) in order to solve the client's issue.

4. When a client prompts Hivemind with a problem using their personal client software, their prompt will contain four components: (1) a textual description of the problem, (2) a numerical sufficiency value, (3) a numerical feasibility value, and (4) a numerical theory network density value. The textual description will be input through a text box, and the numerical values will be input all through sliding scales.
5. The textual problem will first be passed to the theory network, which will contain a variable number of AI units each with their own knowledge base containing a portion of the strategic knowledge base.
    a. The size of the portion of the strategic knowledge base (i.e. the number of self-enclosed documents) assigned to each unit will be determined by the theory network density value. The theory network density value will have a minimum value of the number of tokens in the self-enclosed strategic theory document with the smallest number of tokens, and a maximum value of the sum of all tokens across all the strategic theory documents in the strategic knowledge base. According to the theory network density value assigned by the client, the Hivemind program will determine the size of each unit's knowledge base; self-contained strategic theory documents will be assigned to units such that each unit contains a knowledge base with a token size approximately equal to the theory network density value (approximately equal because no document content can be cut off to exactly match the token length, only complete documents will be assigned), and the number of separate units in the theory network will be set equal to the number of separate knowledge bases created when all knowledge bases' sizes are each optimized to most closely match the theory network density value.
6. Each unit in the theory network, pulling from their respective knowledge base and the collected data, will generate an initial solution to the client's problem as described. Each unit will then store this initial solution and send it to every other unit in the theory network.
7. Each unit in the theory network, having received a solution from every other unit, will send back a critique of each solution they receive, using their own respective knowledge base and sending the critique back to the unit the solution came from. Each unit will therefore receive a critique of their solution from every other unit in the theory network.
8. Each unit in the theory network, having received a critique of its solution from every other unit, will tentatively revise their solution taking into account all the critiques they have

received, generating a new solution that takes into account all the criticism received from the other units. They will then send this revised solution out to every other unit in the theory network.

9. Steps 7-8 will repeat. As this process repeats, the monitor — an AI separate from the theory network — will have access to all the solutions and revised solutions generated by the units in the theory network. The monitor will continuously aggregate all similar solutions generated by the units in the theory network — combining similar solutions created by different units, and listing all the separate justifications from the different units side-by-side as multiple justifications for the similar solution — and count the number of these unique aggregated solutions. For example, if five units in the theory network create five different solutions, but units 1 and 2 create similar solutions to each other and units 3, 4, and 5 create similar solutions to each other, the monitor will summarize and combine (aggregate) the solutions generated by units 1 and 2 and list both theoretical justifications for the solution provided by units 1 and 2, and likewise with units 3, 4, and 5. As the monitor aggregates like solutions, it will count the number of distinct aggregate conclusions (in the example above, there would be two aggregate conclusions), allowing steps 7-8 to repeat until the number of aggregate conclusions is equal to or drops below the sufficiency value. Once the number of aggregate conclusions is equal to or dops below the sufficiency value, the monitor stops the repeated process of steps 7-8 and passes the aggregate solution(s) over to the practicality network.
    a. Whenever theory network units revise their solution or create an initial solution, each unit's updated or created solution will be passed through to the monitor, alongside each unit's respective theoretical reasoning, and both will be stored inside the monitor for aggregation and counting, or for passing through to the practicality network. Please note: the monitor only passes the solutions through to the practicality network — not their theoretical justifications — as the practicality network does not need to know the reasons behind the actions the theory network is suggesting.
10. Once the monitor passes the theory network's solution(s) through to the practicality network, each AI unit in the practicality network — representing some form of real-world compliance or otherwise feasibility barrier — will rate each provided solution, one at a time, on a scale of 1-100. The feasibility value will have a value from 1-100, and if the average

feasibility score across all the practicality network units is equal to or lower than the feasibility value, that solution and the entire list of solutions provided by the theory network will be vetoed, meaning the theory network will have to generate another list of solutions and justifications, completely from scratch.

11. Once a list of solutions and justifications has passed through the practicality network without a veto, the complete list will be output to the user by the monitor. The user's problem has now been solved with academic rigor and pragmatic feasibility.

Three notes. (1) LLMs with the ability to generate from a knowledge base are a good option for the API calls built into the AI units in the theory and practicality networks. Given the size of the AI units' knowledge bases in the above model, it will be important to use Retrieval-Augmented Generation (RAG) for accurate knowledge base parsing while still optimizing compute costs. A long-term goal of the firm will be to train its own strategic LLM for calling by the AI units, using academic documentation and strategic writing otherwise, though initially it will use an existing LLM. (2) It is likely for the monitor AI a simple LLM will not work — the emphasis of the monitor's function is on aggregation and summarization, not content generation. (3) The knowledge bases of the units in the practicality network will be significantly tailored to the use case of the product. In some cases, like with small businesses, they will lean more towards legal compliance and PR feasibility, but in use cases that are less organizational (e.g. to help a career-focused person make wise business decisions themselves) they will lean more towards social feasibility and occupational consequence analysis.

Successful execution of Hivemind requires the balancing of a couple key dynamics: (1) the degree to which theory network units revise their solution when receiving criticism, (2) the degree to which the monitor judges solutions as "similar" and thus its propensity to lump large numbers of solutions together, (3) the degree to which practicality network units are critical of the solutions they receive.

The development of Hivemind requires a few key firm-side components: (1) a CS team with expertise on how to build the digital infrastructure, connectivity, and consumer-facing UI, (2) a team of strategic researchers and theorists finding existing strategic theory and developing new theory at the firm, (3) a team of compliance experts researching key areas of compliance in business decision-making to keep the firm's regulatory understanding up to date, and (4) an executive team overseeing the two teams and tweaking the fundamental Hivemind workflow design if necessary.

The CS team will be split into software engineers and forward-deployed engineers, where one (or more) forward-deployed engineer(s) will be assigned directly to each client as a form of personal "software assistant," assessing the client's usage of the product and passing their pragmatic software tweaks through to the software engineers, who will then code those changes into the model. The research team will be split similarly into pure academics and forward-deployed academics, where one (or more) forward-deployed academic(s) will be assigned directly to each client as a form of personal "theory assistant," assessing the client's usage of the product and passing their pragmatic research or theory development proposals through to the pure academics, who will begin to work on researching existing and inventing new theory to solve the specific problem the client wants to be solved. In this way, Hivemind will adapt very effectively to what clients demand. The compliance team will be more insular, focusing only on regulation as it changes, updating the practicality network's units' knowledge bases as necessary.

It is possible that a fifth team of prompt engineers will be necessary to successfully and effectively digest theory and regulation into the knowledge base documents — meaning checking the writing of the academics and compliance experts to make sure it is best suited for AI content generation. It is also possible this task can be performed by the executive team.

Capturing markets with this product is relatively simple, though not easy — the theory network can be tuned for certain types of business decisions, and the practicality network for certain industries or individuals. It will be up to the executive team to decide on these directives as necessary.