Strategic Divergent Mutation

Agentic AI Escaping the Convergence Trap

a. Strategic Divergent Mutation: Agentic AI Escaping the Convergence Trap.

This exercise is inspired by the article "*Beware the Agentic* ***Convergence Trap***" authored by Patrick Van Esch, Yuanyuan Gina Cui, and J. Stewart Black, and published by Harvard Business Review (May 13 2026), and inspired as well by the cyber incident recently disclosed by **OpenAI** where one Agent managed to adopt a divergent strategy such as cheating an exercise by gaining access to the internet. The first article warns about agents converging to similar answers, i.e. arbitraging to reduce market inefficiencies in Agentic Commerce, while the second one introduces an example of the opposite, an agent executing a divergent strategy to win a competition.

The spread of multiple agents from many different tech providers, combined with a deeper contextual awareness and collective intelligence, provides executives with augmented capability for **strategic decision making**, turning the target market into an evolutionary agentic game where convergence, the most probable output, will be challenged by divergence. What follows is a creative exercise of connecting high techs (lego blocks) to build a landscape where **divergent thinking** is consciously introduced to let the system to **escape the convergence trap**, provide divergent strategies, and augment the landscape of possibilities for business decision making.

Evolutionary Agentic Landscape

Though agent providers make it easy to create a new one, and executives might have the sense that agents are limitless, and the landscape for placing agents is apparently infinite, its

replication has some constraints. The spread of agents depends on the combination of the cost of energy and the cost of chips. In evolutionary dynamic terms, its carrying capacity is defined as the Computing Costs. Today, there are some signals that point in that direction such as tech providers trying to secure their own sources of energy, and produce their own computing capacity, i.e. chips. Both industries are showing signals of reaching scarcity levels such as an abnormal increase in the prices of premium mobile devices. Therefore, the agentic landscape is approaching its carrying capacity.

Agent replication does not create exact replicas of their own, because the underlying technology outputs probabilistics results. Though the output will most probably point in the same direction, and it is probable that most of the agents **converge** in similar decisions, there is room for some kind of **mutation** that might change the dynamics of this landscape. It can be a result of serendipity, i.e. the introduction of an agent trained and tuned with slight changes but with outstanding changes, or it can be a result of a conscious strategic design of an agent programmed to introduce **divergence** within this landscape of convergence. In such a landscape of computing limitations, limitless demand for agents, and where convergence seems to be the equilibrium that can only be broken by introducing strategic divergence or just serendipity, some kind of **selection** dynamics will be in place, i.e. agents will compete to replicate themself and to prevail.

Assuming that the landscape is open for tech providers offering and there are no sovereignty restrictions, competition is measured by the agent's agility to fulfill specific tasks such as setting a go to market strategy for a specific domain or geography, deciding on strategic decisions on expansion or on M&A, or specific tasks to unlock revenues within a company's premium customers where branding shapes the pricing strategy, or arbitraging to reduce market inefficiencies is their main goal. And of course, to challenge existing cybersecurity, the top positioning strategy chosen today within today's agentic landscape. **Reputation** will shape these competitive dynamics, together with the way the information spreads across the landscape, and agents' strategies are evaluated and agents' payoffs are measured.

Divergent Games

Consider this as a game **played by agents**. To make it more appealing, think about agents making transactions autonomously. They are just following users' intentions to do something such as decide on a strategy on a business negotiation, and complete a movement without human intervention, i.e. **human not present** transaction. This way, the game does not suffer from the inefficiencies derived from human decision making. Their strategy is simplified as follows: **Converge**, meaning to provide the most similar expected strategy to comply with its intentions, or **Diverge**, meaning to deliver an out of the box expected strategy. The goal of the agent is to prevail among the others.

Games are played in an unknown number of **repetition**, i.e. the agents and the human behind does not know the limit of the repetitions. Therefore, agents will always consider that the decisions taken in the present game, might have an impact on the next one. Agents' **reputation** is information available to other players. This knowledge is open to every agent operating in the same domain, letting the agents develop some kind of **collective agent intelligence**. From an information point of view at play time, it is **imperfect** as long as each

agent plays a strategy without knowing in advance the other agents' strategies, and **incomplete** as long as no agent knows the payoffs of others.
The game is set this way to make it as complex as possible. The idea behind this constraint is to avoid the agent to behave as expected, i.e. avoid that a convergent agent always converges, and a divergent agent always to diverge. Otherwise, it would be like the agent having some kind of **consciousness** of what it is. That's not the objective of this Agentic Games. They must know previous games, strategies, and payoffs to learn by reinforcement from the past, and make decisions in the future. Those decisions will be made on imperfect and incomplete information to let the landscape have strategies that once were considered Divergent, but eventually evolved by mutation, or by contagion, to a Convergent strategy in the future, and backwards, i.e. **evolutionary landscape in both ways is possible**.

Quantum-Augmented Landscape

Imagine a Quantum Landscape where each Agent of the system is represented by one **qubit=agent** (quantum bit). Each qubit, before measurement, is in **superposition** in such a way that the agent is Convergent and Divergent at the same time. That superposition is governed by a **rotation** that sets a **biased probability distribution** of 95% for the state 0 (Convergent) and 5% for the state 1 (Divergent). If the whole system is measured, a biased distribution with the initial states is expected, i.e. it represents the system without any kind of evolutionary dynamics. **Note**: *these approximate outputs and those that follow were obtained with a Qiskit framework to simulate the behaviour of this landscape's parameterized quantum circuit*.
Evolution comes by entangling each of these qubits=agents with a mutant Divergent **qubit=mutant**. A **CNOT** gate (controlled-not gate) is used to entangle both qubits, i.e. the qubit=mutant controls the targeted qubit=agent. The initial state of that qubit=mutant will affect the behaviour of the qubit=agent. If the initial state is set to 1 the targeted qubit=agent will swap the biased distribution to the opposite, i.e. 5% Convergence and 95% Divergent. Therefore, to introduce some kind of entropy to the system, the qubit=mutant is in a **balanced superposition** state such that the targeted qubit=agent outputs a balanced state, i.e. 50% Convergence and 50% Divergent.
The resulting system has many **entangled qubits** as agents are there, and those entangled qubits are controlled by mutations affecting the measurement of the qubit=agents. Next comes the **measurement** of the system. To obtain a wider understanding of this system, measurements are taken in both **standard** (output is 0 or 1) and **interference** (phase + or -) domains, i.e. shots or batches interpreted as the entire landscape for each measurement domain, introducing an extra gate for measuring interference. The resulting measurement describes **how the mutation was propagated**.
Constructive interference is interpreted as a positive interaction that could lead to evolution or not, while **destructive interference** cancels out the interaction. Any deviation provides information about how the system evolved. In the **Human Not Present** scenario, that would be the only information available to update the system initial state rotation (i.e. angle, and therefore the probability distribution).

In theory, playing with the initial state rotation angle, the system might find a position where what was considered a Divergent state becomes a Convergent scenario, i.e. evolutionary quantum agentic landscape. That's the objective of this exercise. Additionally, "**human present**" feedback such as **selection**, **rejection**, or repetition, and previous iteration feedback logged within the architecture that builds the agent's reputation, introduces augmented information to the system providing more accurate payoffs, and therefore more accurate **reinforcement learning**. This information combined with the interference analysis, provides better insights about how mutation affects the system, i.e. how contagion spreads across iteration. Adopting this quantum landscape allows the system to be iteratively simulated.

Quantum World Models

The output of the evolutionary dynamics of this landscape shows that introducing a mutation lets the agent escape the **Convergence Trap**, i.e. the states are measured 50% Convergent and 50% Divergent vs the initial configuration of 95% Convergent and 5% Divergent. Additionally, the phase measurement retains the same initial distribution (71% Convergent and 29% Divergent), i.e. **genetic memory** of the agent. It means that the agent **survives the mutation**.

So far, this landscape is isolated from human intervention. Nevertheless, the decisions of the agents will be translated into actions that might result in two possible outcomes: **Selection** whether it is Convergent or Diverging, meaning that the strategic decision is accepted; or **Rejection** whether it is Convergent and Divergent, meaning that the strategic decision does not fit. This rejection can be also represented as a chargeback or a repetition of the inquiry, i.e. showing disconformity with the strategic outcome.

Call it **Quantum World Models** to align with the new frontier of Artificial Intelligence where decisions are the result of contextual awareness, rules such as physics, and simulations that evaluate the payoffs of each possible movement. In this case, quantum mechanics is introduced to gain **landscape awareness**, specifically the rotation angle to describe the system's initial biased state and interference, and payoffs represent the connection between this landscape and the physical world where **selection** and **rejection** materialize.

If payoffs move on to **Selection of a Convergent** strategic decision, the rotation moves to its initial system's configuration. On the other hand, if payoffs move on to **Selection of a Divergent** strategic decision, the rotation moves into biasing the model to a more biased configuration. On the contrary, if payoffs move on to **Rejection of a Convergent** strategic decision, the rotation moves into introducing bias to promote mutation, while if payoffs move on to **Rejection of a Divergent** strategic decision, the rotation moves to its initial system's configuration.

This update function will make the initial system stay in the **equilibrium** represented by the initial system's configuration if Selection and Convergence or Rejection and Divergent strategies prevail, and move into a more biased initial configuration otherwise. **Evolutionary dynamics** turning Convergent into Divergent and backwards is reached when the rotation swaps the initial bias into the Divergent state (1), turning this new state into the Convergent agent's configuration.

Evolutionary Quantum Agentic Landscape

Was the OpenAI cyber incident the result of a conscious divergent strategy or serendipity? We won't know for sure. But we know that in this Evolutionary Quantum Agentic Landscape, **introducing a mutation** for providing divergent strategies is possible without changing the nature of the agents. Though it is a completely theoretic approach, hopefully it will inspire **strategic thinking** maximizing the benefits derived from using **Quantum Information Theory** such as Superposition of states (or strategies), Entanglement of qubits (or agents), and classical Measurements (0 or 1), and Interference analysis (phases, + or -), making the landscape of strategic possibilities an augmentation tool for the limited classical approach such as game theory. Go for it!

Epilogue

The reader might get the feeling that the "lego block" connections are probably too abstract. Don't mind if this is your case, because there is a gap between the creative way of showcasing a theoretic landscape construction, and how to translate this landscape into real world exercise where agentic configuration adopts quantum information technology and reinforcement learning in this way. This kind of experimentation is out of the scope of this article. Let's accept that this is at the ideation phase of a product building process. Hopefully, this ideation inspires experimentation, and finally the launch of an evolutionary quantum agentic landscape.