Generative AI Remakes Scenario Planning

Idea In Short

Strategy leaders should stop treating generative AI as a faster way to write planning documents and start treating it as a way to generate more strategic options before committing capital. Coca-Cola used generative models to crowdsource thousands of creative variations on its brand identity, JPMorgan embedded them into investment research to compress days of analysis into hours and Shopify Magic gave small merchants access to ideation tools once reserved for large marketing departments. Each case points to the same recommendation, build a standing practice of generating multiple scenarios and creative variants continuously, then apply human judgment to stress-test and choose among them, rather than commissioning a single AI-written plan once a year. Organizations that adopt this generate-then-curate discipline will out-maneuver rivals still running annual planning cycles built on one or two static scenarios.

Generative artificial intelligence [AI] has moved past the phase where it merely drafts memos or summarizes reports for corporate strategists. Its more consequential contribution shows up earlier in the process, at the point where teams imagine what might happen next and sketch out how to respond to it. Coca-Cola, JPMorgan and Shopify have each tested this capability in different ways over the past three years and their experience points to a specific opportunity that is often overlooked amid broader AI hype. That opportunity is using generative AI not for administrative speed but for widening the range of futures a company is willing to consider and shortening the distance between a raw idea and a tested concept.

Why Scenario Planning Needed a Creative Engine

Scenario planning has existed as a discipline since Royal Dutch Shell popularized it in the 1970s and it has always suffered from the same structural weakness. Building a handful of credible futures takes weeks of workshops, external research and facilitated debate, so most companies settle for two or three scenarios and revisit them once a year at best. Generative AI changes the economics of that exercise because a model can draft dozens of internally consistent scenarios from a prompt, a data set and a set of constraints in a single afternoon. Strategy teams can specify variables such as interest rates, regulatory shifts or shifts in consumer sentiment and the system will produce narrative and quantitative outputs that a human team would otherwise need days to construct. Harvard Business Review has documented how generative models can synthesize signals from news, filings and social data into scenario narratives that planners then refine rather than draft from a blank page.1 The value is not that the machine replaces judgment, but that it removes the drafting bottleneck that used to limit how many futures a company could seriously entertain before choosing a direction.

This matters because the cost of an unexamined scenario is rarely visible until it arrives. A retailer that only modeled a mild slowdown found itself unprepared when a supply shock hit two categories at once and a bank that planned around gradual rate changes struggled when a policy reversal happened within a single quarter. Generative tools let planning teams run many more of these permutations before committing capital and they can do so continuously rather than as an annual ritual. That shift, from a periodic planning event to an ongoing modeling capability, is the real transformation underway in corporate strategy departments right now.

Coca-Cola's Create Real Magic Experiment

Coca-Cola offered one of the earliest visible examples of generative AI applied to creative strategy when it partnered with OpenAI in March 2023 to launch Create Real Magic, a platform built on customized versions of GPT-4 and DALL-E. The company invited digital artists and consumers around the world to remix its brand assets, including its iconic contour bottle and signature script logo, into original artwork submitted through an open platform. Coca-Cola described the initiative as a way to test how generative tools could extend brand storytelling beyond what its internal creative teams could produce alone, turning thousands of outside contributors into an ad hoc ideation studio.2 The submissions were reviewed and select pieces were later featured on digital billboards in New York's Times Square and London's Piccadilly Circus, giving the exercise commercial stakes rather than treating it as an internal pilot.

What makes the case relevant to strategy, rather than marketing alone, is the underlying mechanism. Coca-Cola effectively used a generative model to run thousands of parallel creative experiments against its brand identity at a cost and speed no internal studio could match. Strategy teams evaluating new product concepts, packaging directions or regional campaign ideas can apply the same logic on a smaller scale, generating many variants of a positioning idea and letting quantitative signals, rather than a single creative director's instinct, narrow the field. The lesson from the campaign is less about the specific artwork produced and more about the discipline of using generative volume to surface options that a conventional brainstorm would never reach.

JPMorgan's Shift From Reports to Real-Time Modeling

JPMorgan approached generative AI from a different angle, embedding it directly into the machinery of investment research and portfolio analysis rather than brand ideation. The firm's research division has published extensively on how generative models process earnings transcripts, filings and market commentary to produce draft analysis that human analysts then verify and refine, compressing a process that used to take a full day into a matter of hours.3 Analysts still set the questions worth asking and still sign off on the conclusions, but the volume of source material a single team can synthesize has expanded considerably because the model does the first pass of pattern recognition across datasets too large for manual review.

The strategic implication reaches beyond the trading desk. A planning team weighing a market entry decision can ask a similarly configured model to scan filings, patent activity and hiring trends across dozens of competitors simultaneously, producing a structured summary that a strategist would otherwise need a research team of several analysts working for weeks to compile. JPMorgan's experience suggests the more durable advantage lies not in any single report the AI produces, but in the compounding effect of running that kind of synthesis continuously, so that shifts in a market or a competitor's posture surface within days rather than at the next scheduled review. Strategy functions that treat this capability as a one-time research boost, instead of a standing input to how they monitor the environment, are likely to underuse it.

Shopify Magic and the Democratization of Strategic Creativity

Shopify took the idea of AI assisted ideation and pushed it down to the smallest businesses on its platform through Shopify Magic, a suite of generative tools launched in 2023 that drafts product descriptions, marketing copy, email campaigns and social posts from a short prompt. The company built Magic and its companion assistant, Sidekick, to let a merchant with no marketing staff produce campaign material that would previously have required an agency retainer, effectively giving a single founder access to a creative and strategic planning function that used to belong only to larger organizations.4 A merchant can generate several variations of a product description tuned to different customer segments, compare their tone and framing side by side and select or blend the strongest elements before publishing.

The significance for larger organizations is instructive precisely because Shopify built the tool for the resource-constrained end of the market. If a solo entrepreneur can generate and test a dozen positioning angles for a new product in the time it once took to write one, a corporate strategy team with far greater data access and computing budget has little excuse for sticking to a single narrative when it pressure-tests a new initiative. Shopify's approach also demonstrates that generative ideation works best as an iterative loop rather than a single output, since the tool is designed for a merchant to generate, edit and regenerate rather than accept the first draft outright. That iterative pattern, generate broadly and then curate deliberately, is the same discipline that separates useful scenario planning from a stack of AI-written slides nobody reads.

Building Scenarios Leaders Actually Use

Generating options is only half the exercise and the harder half is turning generative output into scenarios a leadership team will actually use when a decision is on the table. McKinsey has found that the organizations getting the most value from AI in strategy development are the ones that pair generative tools with a disciplined process for stress-testing the assumptions behind each scenario, rather than treating the model's first output as a finished plan.5 That distinction separates companies that use generative AI as a genuine strategic input from those that use it as an expensive way to produce more paperwork.

From Static Slides to Living Models

The traditional scenario deck gets built once, presented to the board and then filed away until the next planning cycle begins. A generative approach turns that deck into something closer to a living model that a strategy team can re-query as conditions change, adjusting a single variable, such as a tariff rate or a shift in raw material cost and immediately seeing how the downstream narrative and financial projections move. This responsiveness matters most in industries where the operating environment turns over faster than the annual planning calendar, including retail, financial services and technology. Teams that adopt this approach report spending less time rebuilding scenario decks from scratch and more time debating which scenario the current evidence actually supports.

Stress-Testing Assumptions Before They Break

A scenario is only as good as the assumptions embedded in it and generative models are well suited to surfacing assumptions a human team might not think to question. Prompting a model to argue against its own baseline scenario, or to identify the three assumptions most likely to be wrong, forces a kind of adversarial review that most planning teams skip because it is time-consuming to run manually. This practice works best when a named strategy lead owns the output and is accountable for deciding which challenged assumption warrants a revised plan, rather than leaving the model's critique to sit unread in an appendix. Organizations that build this adversarial step into their planning cadence tend to catch blind spots earlier than those that treat the first generated scenario as sufficient.

The Human Judgment That AI Cannot Replace

None of this removes the need for a leadership team to decide which future to bet on and that decision still rests on judgment, appetite for risk and knowledge of a company's own capabilities that no model currently possesses. The World Economic Forum has emphasized that organizations succeeding with generative AI treat it as an amplifier of existing strategic capability rather than a replacement for the people who set direction, pairing every AI-generated recommendation with a clear line of human accountability.6 A model can propose that a company enter a new geography or divest a business line, but it cannot weigh the reputational cost of a decision the way an executive who has spent a career inside that industry can.

> A generative model can widen the field of options a leadership team considers, but it cannot decide which risk that team is willing to carry

The most effective strategy functions use generative output as a starting position for debate rather than a conclusion, deliberately assigning someone in the room to argue for the scenario the model rated least likely, since low-probability events are precisely the ones that catch organizations unprepared. This discipline requires strategists who understand both the business and the limitations of the tool well enough to know when a generated scenario reflects a genuine signal and when it merely reflects patterns in the data the model was trained on. Building that judgment takes deliberate practice and companies that skip the practice in favor of speed tend to discover the gap only when a scenario the model rated as unlikely turns out to be the one that materializes.

Governance, Talent and Responsible Adoption

Adopting generative AI for strategic work is as much an organizational project as a technical one and companies that treat it purely as a software rollout tend to see adoption stall after the initial pilot. Strategy teams need practical fluency with prompting, data preparation and output verification, which means training programs that go beyond a single onboarding session and instead build the skill into how analysts do their daily work. Some organizations have started designating strategists as internal prompt specialists, tasked with refining how the rest of the team queries these systems and maintaining a library of prompts that reliably produce useful scenario structures.

Data governance carries equal weight, since a generative model trained or fine-tuned on proprietary market intelligence, customer data or competitive research needs the same access controls and audit trails that any other sensitive corporate asset requires. Leaders should also set explicit boundaries on which decisions require human sign-off before a generated scenario can influence a resource allocation choice, rather than letting that boundary emerge informally after a mistake occurs. Companies that establish these boundaries early tend to move faster over time, not slower, because their teams trust the system enough to rely on it for the ideation and modeling work it does well, while reserving judgment calls for the people who are accountable for them.

- 1Use GenAI to improve scenario planning
- 2Coca-Cola invites digital artists to Create Real Magic using new AI platform
- 3The rise of generative AI
- 4Shopify Magic and Sidekick: AI for commerce
- 5How AI is transforming strategy development
- 65 ways to make the generative AI transformation a success

Summary

Generative AI has proven most valuable in strategy not as a report-writing shortcut but as a way to widen the field of futures and creative options a company considers before it commits resources. Coca-Cola's Create Real Magic platform, JPMorgan's embedded research models and Shopify Magic each show a version of the same pattern, generate broadly and then curate deliberately with human accountability attached to every choice. Scenario planning built on this rhythm becomes a continuous capability rather than an annual event, catching shifts in markets, regulation and consumer behavior while there is still time to respond. The organizations pulling ahead are not the ones with the most sophisticated models, but the ones that paired generative volume with disciplined stress-testing and clear ownership of every decision the technology informs.