

PERSONALIZATION

The Four Gaps in Next-Best Action Programs

By [Silvio Palumbo](#), [Mark Abraham](#), [Karl Johnson](#), and [Yun Lim](#)


ARTICLE MAY 29, 2026 8 MIN READ

This is Part 1 of a five-part series on next-best action.

Next-best action (NBA) is the crown jewel of personalization strategy. The promise is elegant: blend data, technology, and science to determine the right message for the right customer at the right time at a speed and level of precision not possible with manual approaches. Hundreds of millions of dollars have flowed into platforms, data infrastructure, and analytics teams to deliver on that promise. But current implementations have plateaued, and NBA needs to take the next evolutionary step.

Most enterprise NBA programs today represent a compromise between scale and automation. As campaign volumes rise, journey counts multiply. Eventually, operational execution overwhelms personalization teams, often leaving them unable to point to tangible value creation. Because the operating environment evolves faster than the systems built to manage it, four structural gaps—in architecture, science, operating model, and measurement—have widened and reinforced each other. (See the exhibit.)

The Four Gaps That Cause Most NBA Programs to Underdeliver



Source: BCG analysis.
Note: NBA = next-best action.

Closing these gaps and enabling NBA to evolve to the next level will require more than an incremental improvement. Companies need to close all gaps simultaneously, which they can do by embracing technology shifts related to foundational models, composable architectures, and reinforcement learning.

The Architecture Gap: The Journey Model Is Breaking

The average enterprise now manages thousands of customer journeys. Most of those journeys overlap, conflict, or lose relevance within months of launch. Maintenance costs grow linearly while marginal returns plateau. Moreover, the existing structure fails to take into account that customers experience moments, not journeys. A customer calling to resolve a billing issue shouldn't be treated as though the interaction is a winback journey and has no reason to suppose that the company views it as such.

The underlying problem is that decisioning logic can't keep pace with accelerating complexity and the need to manage cross-journey orchestration. Signals, channels, and customer states have grown exponentially, but the decisioning layer operates within a bounded, predefined action

space. The architecture was designed for a world in which humans define the option set and models optimize within it. That paradigm can now evolve.

What comes next will be a shift from marketer-orchestrated journeys to agent-composed actions. Marketers will retain the creative task of populating the universe of possible options from a modular shelf of assets. But AI agents will then select, sequence, and compose personalized interactions from that shelf in real time. This is not a platform upgrade. It is a different unit of decision based on individual actions, not broad journeys. Part 2 of this series lays out the technology that enables this architecture.

The Science Gap: The Models Optimize for the Wrong Thing

Unsophisticated NBA systems rely on propensity models that predict future behaviors such as next purchase, next item, and preferred channel. In many instances, the models are well built and technically sound, but they only answer an intermediate question.

What matters is not predicted behavior, but changed behavior. How much does an action alter the outcome compared to doing nothing? A customer with a 90% propensity to purchase may have planned to make the purchase anyway. Offering a discount to customers who fall into that category destroys margins without changing behavior.

More sophisticated implementations have already shifted from propensity modeling to uplift modeling. During that transition, organizations typically discover that 20% to 40% of their active programs delivered negligible incremental lift. Although they improve on their predecessors, uplift models are slow to recognize novel combinations from an expanding composable shelf, and they treat each decision in isolation rather than as part of a sequential engagement strategy.

Closing the science gap requires a broader decisioning stack consisting of three layers:

- Propensity and uplift models for foundational scoring
- Contextual bandits for real-time, personalized decision making that balances exploring new options with exploiting known high-reward actions
- Foundation model agents for compositional reasoning that is beyond the capability of structured models

Part 3 of this series unpacks the science underlying the new model.

The Operating Model Gap: The Team Is Structured for a World That Is Starting to Disappear

Most marketing organizations are still structured around campaign execution, with teams organized by channel, life-cycle stage, or product line. Each team builds, launches, and manages programs within its domain. Coordination happens through campaign calendars, journey maps, and weekly standup meetings. This structure has evolved around humans as the primary decision makers and around campaigns as the primary unit of work.

With the aid of agent-native systems, organizations can move beyond campaigns and journey flows. The agent will initiate actions in real time from a modular shelf. While the system makes thousands of individually composed decisions per hour, there will be no need for human marketers to conduct weekly standups to discuss which campaigns to execute. Instead, they can ideate the best possible inputs to make available for campaigns.

The marketing team will fundamentally shift from building campaigns to curating the shelf, from defining journeys to setting objectives, and from quality assurance review to continuous governance. We anticipate that the marketing value chain will compress from months to days.

Part 4 of this series examines how the marketing operating model will transform, and what the resulting new roles, skills, and governance structures will look like in practice.

The Measurement Gap: Most Programs Cannot Prove Their Value

Proving ROI from NBA is complex. Many organizations use only activity metrics such as open rates, click-through rates, and conversion rates. But to measure the value generated by marketing actions that would not otherwise have materialized—and thus to accurately calculate ROI for the

marketing campaign—organizations need to design experiments scientifically, embrace a culture of test and learn, and secure legal approvals in regulated environments.

As agentic marketing accelerates, organizations urgently need to build a new level of rigor and trust in measurement science. For example, without a persistent global control group or a synthetic control group (such as a comparable proxy by geography), an organization can't isolate the NBA program's true incremental contribution over the long term. In the absence of tactical-level experimentation, the organization has no way of knowing which specific actions drive lift and which are neutral. And when its measurement architecture is not designed for continuous learning, the organization has to make investment decisions on the basis of metrics that conflate activity with impact.

Imposing this new level of measurement rigor is not easy and often unearths limitations in teaming, design, and execution. [Part 5 of this series](#) discusses the new measurement paradigm—from campaign metrics to incrementality testing to agentic measurement, where the system designs its own experiments, interprets the resulting data, and acts on its findings within governance constraints.

Three Converging Technology Shifts

Technical leaders have recognized the emergence of these four structural gaps. Three technology shifts that are now maturing and converging will allow organizations to close all four gaps simultaneously:

- **Foundation models** enable agents to understand unstructured context, consider novel combinations of actions, and articulate why they made a decision.
- **Composable architecture** permit agents to draw from a modular shelf of atomic assets, compiled communications, and micro-journeys and assemble personalized interactions in real time.
- **Reinforcement learning** allows agents to learn from their own decisions, rather than replaying patterns in historical data, thereby accelerating exploration and optimization.

Harnessing these technologies to address the four structural gaps will yield a foundational advantage, whereas taking a less comprehensive approach will deliver only partial improvements. An organization that upgrades its models but keeps its journey-based architecture

in place can gain efficiency—but only within a fundamentally constrained paradigm. Similarly, an organization that builds an agent architecture but retains a campaign-era measurement model cannot prove that an investment was worth making.

The Compounding Stakes

Agent-native NBA systems are compounding systems: every customer interaction generates data that improves the next decision. As a result, early adopters of agentic systems gain a lead that widens with every touchpoint. Their decisioning gets smarter, their shelf benefits from better testing, their measurement becomes more precise, and their operating model grows more efficient—simultaneously and continuously.

This five-part series maps the transformation needed to close all four structural gaps. It is written for leaders who are responsible for making the investment decisions, designing the architecture, building the teams, and governing the systems that will define the next era of customer engagement and NBA. The gaps are real, but the path to closing them is now clear. Organizations that move first will earn a significant and enduring advantage.

This is Part 1 of a five-part series on the future of next-best action. [Part 2](#) discusses the shift from marketer-orchestrated journeys to agent-composed actions and the technology architecture that enables it. [Part 3](#) unpacks the decisioning science, from propensity models to contextual bandits to foundation model reasoning. [Part 4](#) examines how marketing organizations can restructure around the new operating model. [Part 5](#) addresses measurement: from campaign-level attribution to agentic measurement.

Authors



Silvio Palumbo

Managing Director & Senior
Partner
New York



Mark Abraham

Managing Director & Senior
Partner; Global Leader,
Marketing, Sales & Pricing
Practice
Seattle



Karl Johnson

Managing Director & Partner
Washington, DC



Yun Lim

Associate Director, Marketing
Chicago



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.