Synthetic Respondents as Real-Time Preference Infrastructure for Agentic Commerce

Technical Executive White Paper

Executive Summary

Agentic commerce systems are moving from prototypes to production. These systems make recommendations, allocate budgets, and execute purchases on behalf of users. In many cases, these decisions must occur in seconds or milliseconds.

This creates two non-negotiable system requirements:

1. Preference subsystems must operate at the same time scale as the agent's decision loop.
2. That data must be accurate and actionable for real purchase decisions.

Traditional market research satisfies the second requirement but fails the first. Heuristic or rule-based systems satisfy the first but not the second. Synthetic respondent systems are the only practical approach that satisfies both constraints simultaneously: they operate at agent speed and, as demonstrated across dozens of validation studies, produce results that align with live survey outcomes within typical sampling variance.

The studies in this paper are not notable because of their findings. A traditional survey would produce similar results. What is new is that these results were generated in seconds, inside an agentic workflow, with distributional alignment validated against live benchmarks.

Key Technical Metrics:

Validation: Validated across dozens of studies in three major domains—healthcare, consumer, and social. Synthetic results consistently reproduce directional findings, segment differences, and decision-relevant patterns observed in live surveys.

Latency: Milliseconds to seconds, depending on pipeline design, query complexity, metrics returned, and hardware configuration—operating at the same time scale as agent decision loops.

Cost: Orders of magnitude lower than traditional research. Traditional live studies typically cost tens of thousands of dollars per scenario; synthetic studies incur marginal compute costs and can be pennies per metric or query.

Coverage: Effectively unlimited scenario variants. Synthetic panels support continuous, segment-specific, and counterfactual queries, whereas live studies are typically single-scenario and expensive to repeat.

---

1. The Preference Timing Constraint

Agentic systems must make decisions such as:

- What product to recommend
- How to allocate a budget across categories
- Whether to substitute an out-of-stock item
- When to act autonomously versus request confirmation

Each of these decisions depends on consumer preferences.

Traditional research pipelines:

- Operate outside the agent's decision time scale
- Cannot run inside live decision loops
- Address one scenario at a time
- Produce static, precomputed results rather than live preference queries
- Require full study cycles and fixed budgets per scenario

Agentic systems:

- Encounter thousands of scenario variations per user
- Must respond in milliseconds

- Require preference answers during the decision loop
- Cannot pause for weeks while research is fielded

This creates a fundamental timing incompatibility between agentic decision loops and traditional research pipelines.

Speed alone, however, is not sufficient. Any subsystem operating inside the agent loop must also produce results that reflect real consumer behavior with quantifiable accuracy.

---

2. From Research Projects to Preference Subsystems

Historically, research has been:

- Episodic
- Conducted in advance
- Used for strategy and planning
- Delivered as reports, not real-time data

Agentic systems require something different:

- Continuous preference access
- Scenario-specific insights
- Generated on demand
- Embedded inside the decision loop
- Delivered via API or microservice

In this architecture, preference data becomes an operational subsystem, not a periodic research deliverable.

---

3. Synthetic Respondents as a Real-Time Preference Subsystem

Synthetic respondent systems are AI models trained on large-scale survey data. Given a scenario and preference queries, they generate the expected preferences for a target population matching the current customer.

From an agentic architecture perspective, this is not a research tool. It is a **decision subsystem**.

Traditional research produces:

- Static reports
- Based on predefined scenarios
- Generated outside the decision environment

Synthetic respondent systems instead provide:

- Scenario-specific preference outputs
- For the current customer segment
- Generated inside the agent's decision loop

This transforms research from a periodic activity into an operational component of the agent stack.

In a typical agentic pipeline, context agents assemble the decision state, planning agents generate candidate actions, and the Oracle acts as a preference-simulation layer that evaluates human reactions to those actions. The decision agent then selects the final action based on these simulated outcomes before execution.

For agentic systems, the key advantages are:

- Operation at the same time scale as the agent's decision loop
- Scenario-specific preference queries for the current customer
- Decision-grade outputs aligned with live research findings

- Effectively unlimited scenario and segment coverage

4. Real-Time Agent Preference Studies

To demonstrate the operational model, three preference studies were generated using synthetic respondents within a single agent-driven workflow.

In traditional terms, each of these would represent a separate research project. In an agentic system, they are simply preference queries executed during the decision process.

These studies are not presented as evidence of superiority over live research. The directional findings are consistent with what traditional surveys would likely show. Their purpose is to demonstrate that:

- The agent can query preference logic on demand
- For specific scenarios
- For specific customer segments
- Inside the decision loop

Separate validation work conducted across multiple domains shows consistent alignment between synthetic and live survey outcomes.

Study 1: Recommendation Logic

Question: How should an agent allocate a fixed budget for a child's birthday party?

Synthetic respondents produced consistent allocation patterns:

- Food and drinks received the largest share
- Entertainment was the second priority
- Themed supplies and party favors formed the middle tier
- Decorations were the lowest priority

Segment differences followed expected economic patterns:

- Lower-income households prioritized essentials
- Higher-income households allocated more to experiences

On autonomy:

- The dominant preference was for agents to suggest options and request confirmation
- Full autonomy was preferred by a minority

These results are typical of what a live consumer study might produce. The distinction is that the preference logic was generated inside the agent's decision flow.

Study 2: Exception Handling

Question: How should an agent respond when a preferred product is out of stock?

Findings showed clear consumer preferences for control and quality:

- Some consumers preferred automatic substitution
- Others preferred notification and confirmation
- Many preferred substitutions only from a pre-approved list

When evaluating substitutes:

- Effectiveness was the primary criterion
- Price similarity was secondary

These patterns align with traditional brand-loyalty and substitution research. Here, they are produced as part of the agent's decision process rather than through a separate research cycle.

Study 3: Permission Architecture

Question: When are consumers comfortable with autonomous purchases?

Across multiple product categories:

- The dominant preference was for agents to suggest and confirm
- Full autonomy was acceptable mainly in low-stakes categories

Price thresholds showed:

- A central range where autonomy was acceptable
- A meaningful minority rejecting autonomous purchases entirely

These results mirror established consumer behavior findings. The operational difference is that they are generated as part of a live agent workflow.

Scenario: Cold-Start Preferences

Problem: A new user opens the application with no profile, no purchase history, and no stated preferences. The agent must make its first recommendation. What preference data does it use?

Traditional systems default to global averages or static rules. An agentic system with access to a preference subsystem can do better.

In this scenario, the agent has three ambient signals available before any user interaction:

- **Geolocation** — where the user is physically located
- **Retailer context** — the store or platform they are using
- **Search or browse behavior** — the product category or query they entered

These signals are passed to the preference subsystem as targeting variables. Geolocation implies regional demographics. Retailer type implies income band and shopping orientation. Product category implies purchase context. The subsystem converts these contextual signals into a demographic segment and returns preference distributions for that segment — the same outputs as Studies 1-3, using the same query mechanism, at the same speed.

The agent now has usable preference priors for budget allocation, substitution behavior, and autonomy thresholds — before the user has provided a single explicit input. As the user interacts and their profile develops, the agent can refine these priors with individual data. But the cold-start problem is solved at the first interaction, inside the decision loop, without waiting for profile data that does not yet exist.

5. What the Studies Demonstrate

These studies are operational demonstrations, not accuracy claims.

They do not argue that synthetic respondents are more accurate than live surveys. Instead, they demonstrate that preference research can operate as a subsystem inside the agent's decision loop while producing results consistent with traditional studies.

Across multiple domains—including consumer behavior, healthcare, and public opinion—synthetic systems have been validated against live survey data. These validations consistently show:

- Alignment on major directional findings
- Preservation of segment differences
- Reproduction of decision-relevant patterns

This indicates that synthetic preference subsystems can provide decision-grade outputs suitable for agentic environments.

The key operational properties demonstrated by these studies are:

- Preference logic generated at the same time scale as the agent
- Scenario- and segment-specific queries
- On-demand preference retrieval during live sessions
- Integration directly into decision logic

6. Design Implications for Agentic Systems

Three consistent design principles emerge from the studies:

1. Default to recommendation mode

Users prefer agents that suggest options but defer final decisions. Only 32% want full autonomy for complex purchases. Agents should present recommendations with explanations rather than executing autonomously without confirmation.

Implementation: Set default behavior to "suggest and confirm" across all categories except explicit user opt-in for specific low-stakes categories.

2. Treat budgets as hard constraints

74% consider budget violations a dealbreaker. Never exceed stated budget limits without explicit permission, even if a better product is available at higher cost.

Implementation: Enforce budget as a hard constraint in recommendation engine; flag budget tradeoffs for user review rather than auto-escalating.

3. Prioritize quality over speed in exception handling

83% prioritize product effectiveness over price when substitutions occur. Users tolerate delivery delays better than receiving inferior products.

Implementation: Filter substitution candidates by quality/effectiveness first, then optimize for price within acceptable quality band. Default to "notify and wait" rather than auto-substitute lower-quality alternatives.

---

7. Validation Across Domains

Synthetic respondent systems have been validated across multiple domains by comparing synthetic survey results directly with live respondent data. These validations span consumer behavior, healthcare decision-making, and public opinion studies.

Across these studies, synthetic results consistently reproduce:

- Directional findings on key questions
- Segment-level differences
- Decision-relevant behavioral patterns

For synthetic opinion surveys, KL values below 0.15 are generally considered decision-equivalent, meaning a research team would reach the same practical conclusions as with live data.

External Validation Studies

Study	Domain	Questions	Avg. KL
Commonwealth/Kaiser PCP Survey	Healthcare professionals	73	0.035
AMA Prior Authorization	Healthcare professionals	24	0.036
Physician Sarcopenia Survey	Healthcare professionals	18	0.042
IFIC Food & Health Survey	Consumer	47	0.038
Walmart Retail Rewired 2025	Consumer	32	0.037
Happy Returns Consumer Survey	Consumer	8	0.041

Full validation reports for each study are available at simsurveys.com/papers.html.

Operational Implication

These results indicate that synthetic respondent systems can:

- Operate at the time scale required by agentic decision loops
- Produce preference outputs aligned with live population behavior
- Support decision-grade logic across multiple domains

Because this simulation layer has been repeatedly validated against real-world data, it can function as a governance checkpoint inside agentic decision loops, allowing systems to evaluate human reactions before executing actions at scale.

---

8. Conclusion

Agentic commerce systems make preference-dependent decisions — what to recommend, when to substitute, whether to act or ask. These decisions happen inside the agent's decision loop, in seconds or less. Traditional research cannot operate at this time scale. It was not designed to.

The practical consequence is that most agentic systems today make these decisions without preference data. They rely on heuristics, assumptions, or generic defaults. This is not a temporary gap that will close as research catches up. It is a structural incompatibility between how research has always worked and how agents must operate.

Synthetic respondent systems resolve this incompatibility. They accept a scenario and a target segment, and return preference distributions — at agent speed, for any scenario the agent encounters, with accuracy validated against live survey data across multiple domains.

The three studies and cold-start scenario in this paper are not exhaustive. They demonstrate a pattern: any preference question an agent can formulate, the subsystem can answer, inside the decision loop, without a separate research cycle. Budget allocation, substitution logic, permission thresholds, and segment-level priors for unknown users are four instances of a general capability.

The choice facing agentic system architects is not between synthetic preference data and traditional research. Traditional research cannot participate in the agent's decision loop regardless of budget or speed of execution. The choice is between synthetic preference data and no preference data at all.

About This Research

Document: Technical Executive White Paper **Target Audience:** CTOs, SVPs of Technology Strategy, Heads of Data/AI, AI Engineering Firms **Date:** February 2025 **Authors:**

- Joel Friedman, CEO/Founder, Simsurveys
- Myles Friedman, Chief AI Architect, Simsurveys

About Simsurveys: Simsurveys provides synthetic survey respondent generation and real-time preference infrastructure for agentic commerce systems. The methodology demonstrated in this paper enables on-demand generation of preference data for any definable scenario and customer

segment, validated against live survey benchmarks across dozens of studies in multiple domains.

Technical Resources:

- Full validation reports and methodology: simsurveys.com/papers.html
- Custom validation studies and pilot programs: contact@simsurveys.com

Contact: For technical deep-dives, pilot program discussions, or custom validation studies: simsurveys.com