

Test pricing before customers see it

How Skuel Simference turns live customer
behavior into synthetic experiments for pricing
and promotions

INFORMATION BROCHURE | 2026



The highest-leverage decisions are the hardest to test

Pricing and promotions can move revenue and margin far more than UX changes, yet testing them with live traffic is risky.

Companies spend millions of dollars A/B testing incremental website changes such as button colors, page layouts, and minor UX tweaks that often produce only marginal gains in conversion. Yet pricing and promotions, arguably the most powerful levers of customer behavior, remain incredibly challenging to test. Running live pricing and promotions experiments can expose customers to inconsistent prices and offers, erode trust, dilute brand equity, create operational and regulatory issues, disrupt revenue, and are simply infeasible to scale across large catalogs.

The result is [that](#) the most critical behavior-driving decisions in commerce are left to stale historical data at best, and intuition at worst.

Skuel is a synthetic-experimentation and optimization platform designed to address this challenge. It leverages the customer-simulation capability of advanced large language models (LLMs) to enable e-commerce brands to evaluate pricing and promotion strategies before deploying them in practice.

How does Skuel work?

Live customer session behavior is a treasure trove of incredibly rich and telling signals that reveal what a customer wants and what prevents them from buying

A customer's web session contains an extraordinarily rich set of behavioral signals: how long they dwell on a product, whether they move back and forth between alternatives, which parts of a page capture their attention, what additional information they seek, and where they hesitate or disengage. Together, these signals provide a detailed picture of what the customer is looking for and what is driving their decision.

Traditional machine-learning methods struggle to interpret this complex, sequential, and highly contextual data. Advanced LLMs, however, are uniquely capable of reasoning over it. They can infer customer intent, identify the factors shaping a purchase decision, and simulate how that customer is likely to respond to a new intervention, such as a price change or promotion. They can also account for changing market context, such as social media trends, news, events, and competitor moves.

Skuel's Simference Engine is designed to harness this capability. Our research has shown that by leveraging these behavioral signals, LLMs can achieve predictive performance on purchase outcomes comparable to bespoke machine learning models trained on aggregate session features over large amounts of data.

Powered by frontier AI research

Founded by career academics and resulting from over two years of research into leveraging LLMs for high-fidelity customer simulation

The reasoning and simulation capabilities of LLMs for consumer behavior prediction are particularly valuable in scenarios where data is scarce or unavailable, such as predicting counterfactual outcomes under alternative, never-tested-before pricing and promotion strategies. This is precisely where traditional machine learning approaches struggle.

However, leveraging this capability reliably and at scale is not straightforward. Even the most advanced reasoning LLMs must be carefully calibrated to become accurate causal customer simulators. And to make large-scale synthetic experimentation economically viable, those capabilities must be distilled into smaller, faster, and more cost-efficient models without sacrificing simulation quality.

Skuel's Simference Engine is the result of more than two years of frontier research in LLM-powered customer simulation by career AI researchers. It combines accuracy, scalability, and cost efficiency to make synthetic pricing and promotion experiments practical for real-world deployment.

From simulation to optimization

Merchandisers and planning teams must decide not only how a promotion will perform, but which products to discount, by how much, and under what conditions.

For luxury brands, the challenge is identifying a curated set of high-impact offers without diluting brand equity or training customers to wait for discounts. More promotion-dependent retailers face the opposite problem: large catalogs and countless possible offer combinations often push teams toward broad sitewide discounts, inherited playbooks, or intuition.

Skuel turns this into an optimization problem. Its engine can evaluate large numbers of candidate promotion portfolios while accounting for objectives and constraints such as margin, inventory, category priorities, discount limits, campaign budgets, and brand guidelines. Planners can specify these requirements in natural language, and Skuel recommends the promotion portfolio most likely to achieve the desired business outcome.

The result is a shift from blanket sitewide discounts or manually designing a handful of promotions and hoping they work to systematically identifying the best strategy within the realities of the business.

Case study: predicting conversion from behavioral signals

To test the core premise behind Skuel's Simference Engine, we conducted an experiment using a large-scale dataset of anonymized browsing behavior from a major European e-commerce platform. The dataset contained approximately 36 million browsing events across millions of customer sessions, including product views, dwell times, searches, product comparisons, add-to-cart events, and purchases.

The study evaluated whether advanced reasoning LLMs and Skuel's proprietary Simference engine could predict purchase rates using behavioral signals alone. The models received no product names, descriptions, images, or other semantic information. Their predictions were based exclusively on pre-purchase behavior: how customers explored products, how long they engaged with them, which alternatives they considered, and how they navigated through each session.

For each customer session, the LLM estimated the probability of purchase. These predictions were then aggregated across all sessions involving a given SKU to estimate that product's overall purchase rate. These estimates were then compared against reality.

The LLMs received no task-specific model training. Their performance was compared with supervised machine-learning models trained on approximately 12,800 SKUs. All approaches were evaluated on the same held-out set of 100 products, enabling a consistent comparison across methods.

Key findings

~95%

performance of the best ML model

75%

reduction in purchase-rate error
compared to Claude Fable 5

73%

reduction in purchase-rate error
compared to GPT 5.5

Skuel Simference matches supervised ML performance despite no training

Skuel Simference achieves a purchase-rate mean absolute error (MAE) of 0.0070, matching the Deep Neural Network (DNN) benchmark and coming close to Gradient Boosting at 0.0066, even though both supervised models were trained on approximately 12,800 SKUs. It also matches their ability to rank products by purchase propensity, slightly exceeding the DNN's rank correlation (0.429 vs. 0.426) and achieving 95% of Gradient Boosting's performance.

Takeaway: Skuel Simference can make behavior predictions in data-poor scenarios with accuracy at par with what advanced ML methods can achieve with large amounts of data.

Skuel Simference substantially outperforms GPT 5.5 and Claude Fable 5

Skuel Simference combines an advanced reasoning model with a proprietary calibration methodology designed specifically for customer-behavior prediction. The impact is substantial.

Compared with off-the-shelf prediction using OpenAI’s GPT-5.5, Skuel Simference reduces purchase-rate prediction error by 73%, lowering MAE from 0.0257 to 0.0070. It also improves purchase-rate ranking by 63%, increasing Spearman correlation from 0.263 to 0.429. Compared with Claude Fable 5, it reduces prediction error by 75%, lowering MAE from 0.0278 to 0.0070. Its purchase-rate rank correlation is also 2.4 times higher, increasing from 0.177 to 0.429.

Takeaway: Access to a frontier LLM alone is not enough: Skuel’s calibration is key to turning advanced reasoning capabilities into accurate and deployable customer simulations.




Figure 1: Mean absolute error in purchase rate prediction across methods




Figure 2: Spearman rank correlation with true ranking of products by purchase rate

Financial impact

The case study establishes that Skuel Simference, without task-specific training, can predict and rank purchase behavior with performance close to trained machine-learning benchmarks. The commercial value comes from applying that capability to predict purchase behavior under alternative pricing and promotion strategies, based on observed session behavior under incumbent strategies. This counterfactual evaluation is precisely where ML models struggle due to a lack of data.

The financial opportunity extends beyond improved prediction accuracy. It can include:

- Less margin leakage from unnecessary discounting
- More productive allocation of promotional spend
- Better alignment between promotions and inventory priorities
- Faster planning cycles and broader scenario exploration

Financial impact can be measured against partner-specific baselines such as incremental contribution margin, revenue, promotion efficiency, inventory sell-through, and planning time.

Design partnership program

Skuel is looking to work with a select few design partners across major consumer sectors to shape the next phase of the platform.

Skuel brings an advanced technological capability that doesn't currently exist in the market, and that can transform pricing and promotions planning across digital commerce. The purpose of the design partnership program is to understand the full scope of pricing and promotion planning and optimization requirements and to determine how Skuel can have a meaningful impact on the planning workflow beyond producing improved predictions.

Working with merchandising, pricing, promotion, finance, data, and technology teams, we will:

- Map how plans are created, reviewed, approved, and deployed
- Identify the most valuable decisions, bottlenecks, and planning constraints
- Define the objectives, guardrails, and trade-offs the optimization engine must support
- Co-design scenario comparison, recommendation, explanation, and approval workflows
- Determine how Skuel should integrate with existing data and planning systems
- Measure both commercial impact and improvements in planning speed and effectiveness

Design-partner input will directly influence Skuel's product direction, workflow design, integration priorities, and optimization capabilities. We are seeking a small number of forward-looking partners across sectors such as fashion and apparel, luxury, grocery, beauty, specialty retail, and general merchandise. To learn more, contact us at contact@skuel.ai

Summary

Skuel turns anonymized, live customer behavior into synthetic experiments for pricing and promotions. It enables teams to evaluate never-before-tested strategies, compare alternatives at scale, and identify the strongest actions within their commercial objectives and constraints.

The result is a faster and more rigorous planning process: teams can explore more possibilities, understand trade-offs earlier, and make better-informed decisions before customers are exposed. Planners remain in control of the objectives, guardrails, approvals, and final deployment.

About the team

Skuel is led by Vijay Kamble, CEO, and Varun Gupta, CTO - academic researchers in causal inference, reinforcement learning, optimization, and large-scale e-commerce systems with experience spanning Amazon, Microsoft, Google, Bell Labs, and leading universities.



Vijay Kamble is an Associate Professor of Business Analytics at the University of Illinois Chicago whose research focuses on applying reinforcement learning and causal inference methods to optimize retail and e-commerce systems. He obtained his PhD in Computer Science from the University of California, Berkeley in 2015. He previously spent a year as a visiting scholar in Amazon's Pricing and Promotions organization, where he redesigned promotion sourcing algorithms for large-scale sales events such as Prime Day. At Amazon, he saw firsthand how counterfactual evaluation based on observational data is a central challenge in promotions analytics. He co-founded Skuel with Varun when their research showed that AI synthetic data can effectively address this challenge.



Varun Gupta is a Professor of Operations and Information Systems at the University of Utah, and an expert in optimization and machine learning for large-scale e-commerce platforms and digital marketplaces. He obtained his PhD in Computer Science from Carnegie Mellon University in 2010. His research focuses on developing algorithms for complex decision systems operating at scale. At Skuel, Varun leads the company's scientific and technological efforts, advancing the frontier of causal customer simulation from behavioral data. He has spent time in various roles at Microsoft, Google, Bell Labs, and was most recently on the faculty of the University of Chicago Booth School of Business.

Frequently asked questions

01 What can Skuel test?

Skuel is primarily designed to test alternative prices and promotion portfolios before they are exposed to customers. We are expanding to testing other types of interventions such as changes in shipping speed and cost, product bundling, etc.

02 What data does Skuel use?

Skuel primarily uses anonymized pre-purchase session behavior, such as product views, hover patterns, dwell time, searches, comparison paths, add-to-cart events, and purchase outcomes under incumbent prices and offers. It can additionally seamlessly integrate market context, such as social media trends, news, events, competitor moves, etc.

03 Can it scale across large catalogs?

Skuel is designed to evaluate large numbers of candidate promotion portfolios, where live testing would be impractical. One of the key technological innovations in Skuel is its ability to run simulations at a quality at par or higher than the most advanced reasoning LLMs at scale without incurring the corresponding high inference costs of these models.

04 Can it optimize multiple goals and constraints?

Yes. Teams can specify outcomes and guardrails in natural language, including margin, inventory, category priorities, discount limits, campaign budgets, and brand guidelines. Skuel searches for the strongest portfolio within those realities.

05 How is Skuel different from conventional forecasting models?

Conventional forecasting models rely on historical outcomes and aggregated features, so they struggle with strategies that have never been deployed and may miss rapid shifts in customer behavior. Skuel reasons over live behavioral signals while incorporating changing market context, helping teams test new strategies against current customer needs and frictions.

06 What evidence supports the approach?

Skuel's accuracy has been extensively tested on e-commerce purchase data. A recent case study on a European e-commerce platform data found that Skuel Simference matched 95% of predictive performance of advanced ML methods without task-specific training. It is also substantially more accurate (~73% reduction in error) compared to naïve customer simulation using off-the-shelf advanced reasoning models such as GPT 5.5 and Claude Fable 5.

07 What systems does Skuel connect to?

Skuel can connect via API or data feeds to systems that capture or store web-session data, including e-commerce platforms such as Shopify, analytics tools such as Google Analytics and Adobe Analytics, CDPs such as Segment, data warehouses such as Snowflake and BigQuery. When session data is not stored or API access is unavailable, Skuel can deploy its own lightweight data-capture module.

08 How are privacy, security, and human control handled?

Skuel can be deployed on-premises or within a client-controlled VPC, keeping data inside the client's security environment. It uses fully anonymized customer session-behavior data and can be limited to the minimum signals required. Access, retention, and security controls are configured to the client's requirements.