



Automation of Product Decision-Making Based on A/B Testing

Alexander Blinov

Chief Product Officer at Zendrop West Palm Beach, FL, USA

OPEN ACCESS

SUBMITTED 25 June 2025

ACCEPTED 30 June 2025

PUBLISHED 05 July 2025

VOLUME Vol.07 Issue 07 2025

CITATION

Alexander Blinov. (2025). Automation of Product Decision-Making Based on A/B Testing. The American Journal of Engineering and Technology, 7(07), 24–32. <https://doi.org/10.37547/tajet/Volume07Issue07-04>

COPYRIGHT

© 2025 Original content from this work may be used under the terms of the creative commons attributes 4.0 License.

Abstract: This article covers the issue of automating product decisions from A/B tests, trying to knit together what have pretty much been disparate and sometimes even ad hoc stages of experimentation into a single, reproducible, scalable pipeline that includes hypothesis planning, traffic control, streaming analytics, statistical evaluation, and safe rollout. The growth of this inquiry is motivated by the rapid increase in numbers of digital experiments and correspondingly strong demand for A/B testing tools--and the tremendous weakness of traditional manual processes: more than 90% of spreadsheets have errors and one typo in Excel can cost billions undermining the product teams' confidence in the experimental results. The novelty of the work lies in a comprehensive analysis of modern experiment factory architectures that integrate feature flags, Apache Kafka-based streaming analytics, frequentist and Bayesian evaluation methods, multi-armed bandit algorithms, reinforcement learning, and elements of causal ML. A six-layer pipeline concept was proposed in which each stage (from the hypothesis catalog to automatic rollback and result archiving) is implemented by automated means without analyst involvement. Results show that automated A/B processes shrink the experiment cycle from weeks to hours, allow for parallel launch of hundreds of tests, reduce error risk, and speed delivery of winning variants to production. Sequential analysis keeps the false-positive rate under control below 5% along with false discovery rate control; Bayesian modes provide for proper decisions in small samples; and multi-armed bandits plus reinforcement learning virtually eliminate traffic loss during simultaneous exploration and exploitation. The automated system increases the frequency of releases, further improves conversions, and helps improve data-driven culture within

organizations. The paper will be helpful to product managers, data analysts, DevOps engineers, and CTOs who are responsible for building and scaling an experimentation platform and establishing a seamless cycle of product decision-making.

Keywords: A/B testing, experiment automation, streaming analytics, feature flags, sequential analysis, Bayesian methods, multi-armed bandits, reinforcement learning, data-driven.

INTRODUCTION

A/B testing is a controlled experiment in which traffic is randomly split between control and experimental versions of a feature or interface, and the difference in target metrics is interpreted as the causal effect of the change. For digital products, this methodology has become the primary mechanism for hypothesis validation: market leaders such as Google and Microsoft conduct over 10,000 such experiments annually, integrating continuous testing into their development pipelines [1]. Studies indicate that companies systematically employing the A/B approach achieve, on average, up to a 25% increase in conversion, thereby accelerating revenue growth without a proportional rise in costs [2]. It is unsurprising that the global market for A/B testing tools is estimated at USD 850.2 million and is growing at 14% per year, reflecting the demand for data-driven solutions [3].

However, the traditional manual experiment cycle relies on spreadsheets, one-off scripts, and uncoordinated metric calculation methods. Such an approach poorly scales the number of hypotheses that grow, so do delays between ideation and conclusion, and the risk of human error. Audits of financial and product reports reveal inaccuracies in more than 90% of spreadsheets [4], and a historical example from JPMorgan demonstrated that a single Excel typo can cost USD 6 billion [5]. A trust factor also emerges: if the product team questions data correctness, decisions are postponed or made by feel, negating the experiment's value.

Automation aims to eliminate these barriers by converting experiments into a streaming, reproducible procedure, wherein traffic allocation, statistical-power checks, and test termination are executed by a service without analyst intervention. Practical cases demonstrate the scale effect: after adopting automation tools, one e-commerce company reduced the full test

cycle from three weeks to 14 hours, shifting from dozens to hundreds of parallel checks [6]. Shortened decision-making time, reduced error likelihood, and increased release frequency directly translate into revenue growth and competitive advantage, explaining the rapid expansion of automated A/B testing platforms in the market.

MATERIALS AND METHODOLOGY

This study is based on a systematic review of 30 sources, including industry reports, academic publications, and practical case studies. Market size and dynamics for A/B testing were assessed using data from the 9cv9 Career Blog [2], Stanford University [1], and VWO [3]. Research by NextProcess [4], P. Barnhurst [5], and Zobel [9] analyzed the reliability of traditional calculations.

The theoretical foundation encompasses a comparison of frequentist and Bayesian methods for evaluating A/B outcomes: Optimizely Stats Engine describes statistical-significance control and false discovery rate management [18, 20], and the Bayesian mode in Statsig has demonstrated reliable conclusions on small samples [17, 21]. The analysis of streaming analytics architectures is grounded in use cases of Kafka at LinkedIn and Confluent [15, 16], while feature-flag practices are examined through LaunchDarkly reports [11, 14].

Empirical data were collected from case studies at Provar [6], Netflix [22], Uber [23], Airbnb [24], and Booking.com [28]. Additionally, safe rollout and rollback practices [19], LaunchDarkly guidelines [29, 30], and McKinsey's recommendations for formalizing a data-driven culture [12] were incorporated.

Methodologically, a mixed scheme was applied, combining:

1. A literature review of market and technical trends;
2. Comparative analysis of frequentist vs. Bayesian algorithms and multi-armed bandit strategies;
3. Case studies of implementations and their outcomes at Provar [6], Netflix [22], Uber [23], Airbnb [24], and Booking.com [28];
4. Architectural analysis of streaming analytics systems and feature flags;

5. Synthesis of safe rollout practices and error-control measures.

RESULTS AND DISCUSSION

The first controlled experiments originated as early as the 1920s; however, the shift to the online environment did not occur until the late 1990s, when Amazon, Google, and other digital pioneers began splitting traffic between interface variants and tracking metrics in real time [7]. Concurrently, the manual era offers a cautionary tale: according to a recent study, 94% of business spreadsheets contain errors [8], and a single incorrect formula in JPMorgan's Value-at-Risk model resulted in a loss of approximately USD 6 billion—examples that underscore the limitations of manual calculations and drive the pursuit of more robust data-analysis practices [9].

The technological foundation for this transition has been stream analytics: distributed event buses such as Apache Kafka enable the capture of user actions with millisecond latency, thereby allowing metrics to be recalculated and tests to be terminated precisely when sufficient statistical power is reached. It is this capability that LinkedIn, Uber, and Netflix leverage to build real-time solutions atop Kafka [10]. At the next level, feature-flag systems emerge to decouple code releases from feature activation: according to a LaunchDarkly report, 60% of teams began adopting flags within the past year, and 81% have already standardized the practice at least across one or several teams, dramatically increasing release speed and safety [11]. In combination, stream analytics and feature flags form an experimental factory in which a single automated pipeline unites ideas, code, and statistics.

Such a factory is effective only where a data-driven decision-making culture is entrenched. McKinsey research demonstrates that organizations

systematically investing in a data-driven approach improve customer-experience metrics more rapidly and free employees to focus on creative tasks, as routine decisions become automated through analytical services. Recognizing experimentation as the primary source of truth enforces the discipline of hypothesis formulation, metric transparency, and the readiness to roll back unsuccessful changes without political friction, scaling product evolution to an industrial level [12].

The modern experimental factory comprises a service chain that transforms the journey from hypothesis to a validated decision into a fully automated conveyor. The platform's logic is structured around six interconnected layers, each handing off data and control to the next without manual team intervention. The process begins with planning: every growth-related assumption is recorded in a catalog that stores the hypothesis statement, target metrics, preliminary power calculation, and decision status. Experience shows that a unified repository significantly enhances knowledge reuse: in the Eppo Knowledge Base, for example, all completed experiments are indexed for subsequent meta-analysis, facilitating the discovery of analogous ideas and reducing duplicated efforts [13]. Such a catalog is often integrated with Jira or a product-management system, so the transition from idea to launch is enacted as a card-status change rather than by file transfers between departments.

Next comes the traffic-splitting layer. The feature-flag mechanism separates code deployment from feature activation and enables independent services to make deterministic decisions about exposing a variant to a specific user. The approach's popularity is confirmed by the LaunchDarkly survey: 60% of IT leaders consider flags a high priority, as shown in Figure 1, and 79% of companies plan to increase their implementation budget over the coming year [14].




Fig. 1. How do leaders at your company view feature management as an investment? [14]

To ensure instantaneous feedback, user-action events are ingested into a streaming bus—Kafka has become the de facto standard in this role. At LinkedIn, such a pipeline processes over 1.4 trillion messages daily on 1,400 brokers [15], and its operational algorithm is illustrated in Figure 2.

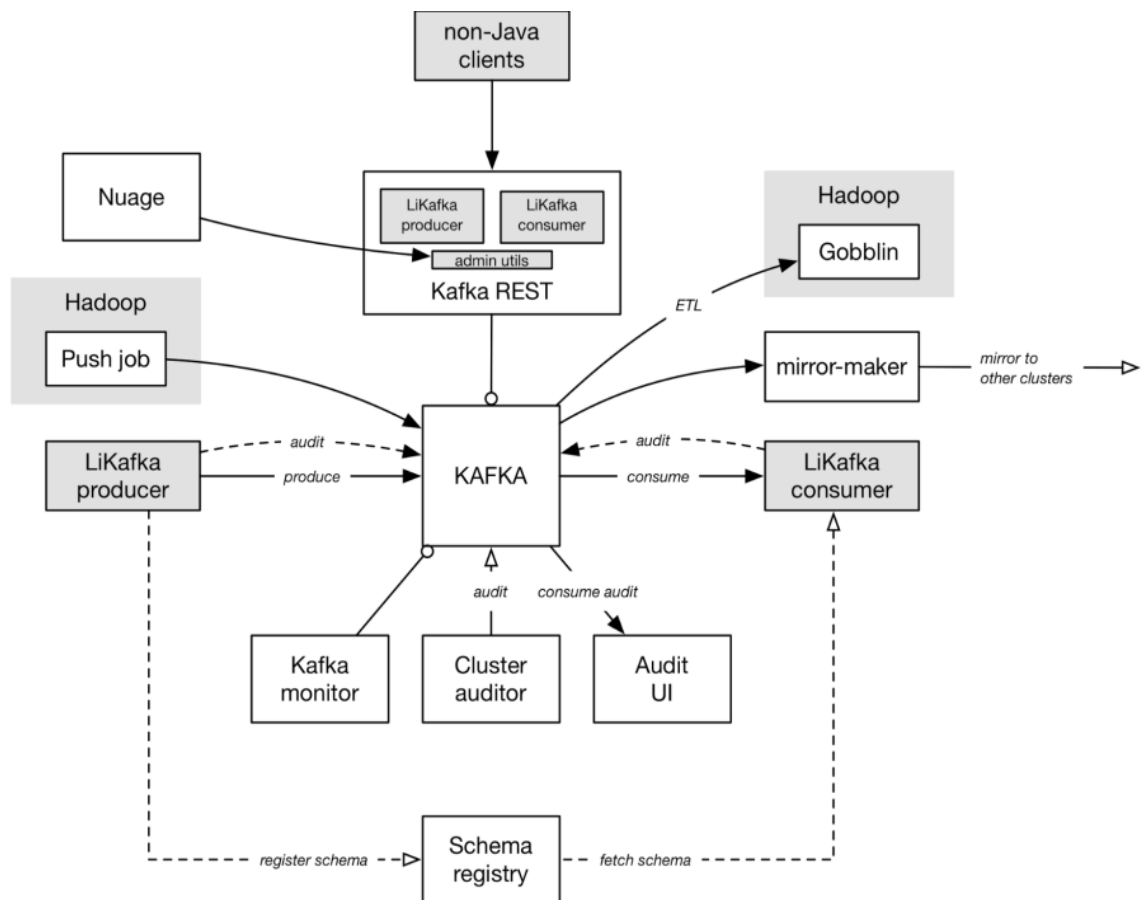


Fig. 2. Apache Kafka Event Streaming and Data Integration Architecture [15]

A Confluent report illustrates the scale of the trend: 79% of the 4,110 IT leaders surveyed consider streaming platforms critical for business agility, and 86% rank them among the top investment priorities for 2024 [16]. At the stream's output, data is aggregated in an analytical warehouse such as ClickHouse, Snowflake, or BigQuery, where incremental data marts are maintained for each metric.

The next layer — the statistical engine — provides basic procedures including the two-sample t-test and χ^2 -test for binary and numerical metrics, but by default allows switching to Bayesian inference, which returns the probability that one variant outperforms another and naturally supports sequential analysis. Study [17] reports that the Bayesian mode enables running tests with samples in the thousands, rather than millions of users, without loss of power, which is particularly important for startups with limited traffic. For loss-optimization tasks, multi-armed bandits are layered atop the Bayesian estimate — they are most appropriate where the delay between action and target metric is minimal. Error-rate control is implemented according to the Optimizely Stats Engine. Until a winner is declared, the system enforces minimum thresholds for visits and conversions, automatically calculating statistical significance and thus reducing the risk of false-positive conclusions [18].

Once the power condition is met, the decision layer engages. Auto-stop terminates the experiment, records the outcome, and forwards it to the auto-roll-out component. The winning variant is deployed according to a predefined rollout curve, with the capability for immediate rollback if metrics degrade. Optimizely describes this scheme as the standard rollout and rollback, permitting flag toggling without recompiling

code and without affecting other tests [19]. The same service generates the events required for incident monitoring. If a degradation beyond the confidence bounds is detected during ramp-up, the system reverts the feature to its original state and opens an alert ticket.

The cycle concludes with the knowledge base and reporting. The engine automatically generates an experiment card summarizing the effect size, p-value, or posterior distributions, screenshots of changes, and segment analysis. A link to this card is attached to the original hypothesis, closing the loop hypothesis → decision → archive. Upon revisiting the same idea, the product team sees the factual history and can decide whether to rerun the test, apply segment-based personalization, or fully archive the hypothesis. Thus, each iteration enriches collective memory and reduces intuition-based decisions, transforming experimentation into a reproducible engineering process.

The experimental factory algorithmic core addresses two orthogonal tasks — reliably measuring effect and minimizing the time between observation and action. At the lowest level remain classical frequentist tests: the two-sample t-test for numerical metrics and the χ^2 -test for binary outcomes. Their strength lies in simplicity and reproducibility, but within event-stream traffic, they easily break down: under continuous peeking, the false-positive rate rises from the target 5% to 57%, while even infrequent checks, once per 1,000 visits, maintain the risk at 20% [20]. An example of the Optimizely interface is shown in Figure 3. Therefore, modern engines wrap the t-test with specialized procedures — sequential analysis and False Discovery Rate control — to guarantee an error rate below the specified threshold, regardless of when the analyst inspects the dashboard.

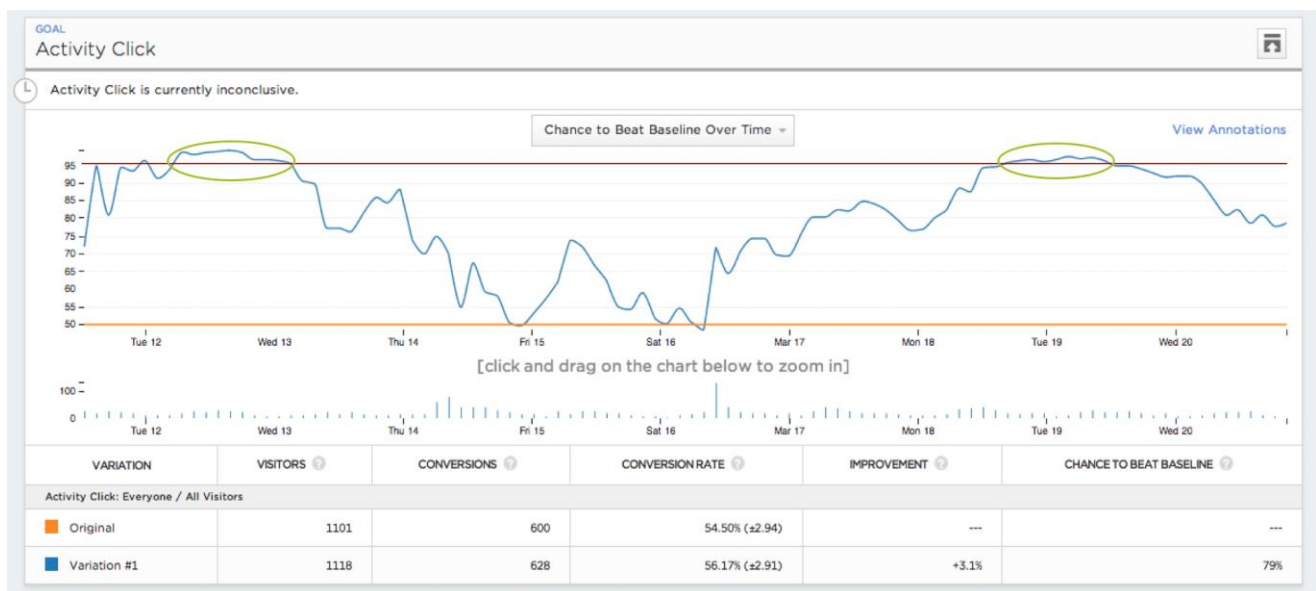


Fig. 3. Interface of Optimizely app [20]

The second layer introduces a Bayesian interpretation: rather than testing a null hypothesis, the system immediately estimates the posterior probability that the variant outperforms the control and the expected loss in the event of an error. Practice has shown that this is particularly useful under low traffic conditions or when multiple metrics are involved: in Statsig, for example, the Bayesian mode permits experiments to be stopped after thousands, rather than millions of observations, while maintaining the risk level at a predetermined threshold [21]. Such acceleration is critical for features whose value rapidly depreciates and for niche segments where assembling a large sample is impractical.

When the objective extends beyond merely measuring the effect to dynamically reallocating traffic, multi-armed bandits and related reinforcement-learning (RL) approaches come into play. Netflix employs contextual bandits to select movie thumbnails and trailers for each user within minutes, reducing the time to identify a winner compared to traditional A/B testing [22]. Algorithms such as Thompson Sampling or Upper Confidence Bound (UCB) initially allocate variants uniformly, then exponentially bias the stream toward the highest-performing option, conserving traffic while concurrently learning. In the streaming architecture described in the previous section, the bandit policy is updated on the fly, and the feature flag adjusts exposure percentages without any code release.

Large-scale ecosystems augment this with an RL component, which optimizes the policy over multi-meter event streams. Uber, for instance, migrated effect

estimation into a client library and achieved a 100-fold reduction in p99 latency (from 10 ms to 100 μ s), enabling parameter adaptation on the order of microseconds [23]. Such speed paves the way for always-on experimentation, in which the system autonomously balances exploitation and exploration.

Finally, atop these layers resides Causal Machine Learning (Causal ML), which addresses whether version B outperforms version A, but also for whom and why?. In practice, this enables immediate post-test audience segmentation: only those clusters with a positive expected effect receive the winning variant, while others continue to see the control or an alternative feature. Thus, the sequence classical \rightarrow Bayesian \rightarrow bandit/RL \rightarrow causal progressively reduces uncertainty, accelerates decision-making, and shifts the success metric from averaged conversion to individualized user value.

Robust implementation of the experimental factory begins with the standardization of metric definitions. As long as different teams debate the definition of an active user, no algorithm can resolve logical ambiguities. Airbnb addressed this issue by developing the Minerva platform, which today encompasses over 12,000 certified metrics and 4,000 dimensions—and it is this unified vocabulary that enables metrics to be transferred instantly from dashboards into A/B reports without rewriting formulas [24]. When planning each test, the catalog immediately records the minimum detectable effect (MDE). For a baseline conversion rate of 15% and a required lift verification of $\geq 10\%$ at 95%

confidence, the system will need approximately 8,000 sessions per variant; a smaller MDE results in slower experiments and more expensive opportunity costs. Such a calculation provides a transparent cost threshold for each idea before its launch and organizes the hypothesis queue according to ROI-to-wait-time ratios [25].

Subsequently, the auto-stop is engaged. Streaming platforms update metrics every minute, and without adjustment, continuous peeking inflates the false-positive rate from 5% to 30% after one week of observations. The sequential mSPRT engine implemented in Statsig maintains an FPR < 5% even with daily peeks; across 560 live experiments, it enabled 58% of tests to be stopped half as early as their planned horizon, while keeping the error probability at 0.4% versus > 20% for a naïve z-test with peaks. In practice, this translates into a precise algorithm: the system evaluates power every N visits, automatically terminates the test upon reaching the boundary, and publishes the decision without analyst intervention [26].

Hypotheses are versioned in a repository akin to Eppo to prevent the idea pipeline from self-blocking. Each entry has a unique ID, a link to the code branch, and automatically inherits the results of previous iterations. Such a ledger forms the basis for a collision detector: when attempting to launch a new test, the engine verifies that no active experiments target the same cohort and metric [27]. The scale at Booking.com illustrates the necessity of this mechanism: the company maintains over 1,000 parallel A/B tests in active memory, and without centralized intersection checks, even an experimental decision risks becoming non-replicable [28].

The final layer is the safe rollout. A feature deemed successful is deployed incrementally, with the kill switch remaining active for as long as the feature itself. Report [29] demonstrates that gradual rollout with immediate rollback capability reduces incidents and accelerates mean time to resolution (MTTR), and [30] records that 76% of elite teams employ this approach. The principle is simple: the flag allows reverting to the control variant in one click until stability is confirmed via monitoring; if a rollback happens, the hypothesis automatically gets a new minor version in the catalog, closing the loop idea → test → conclusion → safe deployment.

The product decision process is automated with A/B

testing, making it possible for the product and marketing strategies to be developed efficiently. Systems that put A/B tests as part of their workflows make decisions grounded on better data, cut down on time for testing of hypotheses, and enhance the total experience of users. Such technologies lessen errors in subjective decisions and open new doors for even more detailed analysis regarding audience needs. In the years to come, further advancements in analytics tools and machine learning should enable much higher degrees of automation and even more personalization, whereby the decision-making process becomes exceedingly adaptable and dynamic.

CONCLUSION

Automated A/B testing helps companies automate the product decision-making since decisions will be made more accurately and faster. It has been proven that switching from the old manual system to an automated one will eradicate major problems plaguing the key issues: human errors, lag delays, and inadequate scalability. Automation tools accelerate the process of testing hypotheses and drastically reduce the time needed to obtain statistically significant results. This, in turn, increases release frequency and reduces costs, which eventually impact revenue growth and reinforce competitive positioning.

Data from automated A/B tests is highly reliable and objective. Automation kills subjectivity in decision-making; it improves analytical processes and makes them more transparent. It also helps in even deeper and more accurate analysis of user needs, which is critical for later personalization of products and services. Using Bayesian methods, multi-armed bandits, and machine-learning tricks on streaming data opens new doors for improving experimental work and its performance when dealing with low—to medium amounts of traffic.

Automated A/B-testing systems set up an experimental factory that turns the making and checking of new features into a steady, shared pipeline, cutting down the time between idea and result. Because of better ways of dealing with numbers, this setup helps firms change products more quickly to suit changing times and reduce dangers linked to wrong answers. Ultimately, these systems become key parts of wider data-based decision-making models that help businesses thrive in today's digital landscape.

Hence, investing in automated A/B-testing tools is a key step toward creating a more capable, effective, and scalable product-making setup. Looking ahead, continued advances in analytics and machine learning are expected to render these systems even more powerful and personalized, unveiling new opportunities for enhancing customer experience and optimizing business processes.

REFERENCES

1. K. Gilbert, "A/B Testing Gets an Upgrade for the Digital Age," *Stanford University*, Jun. 12, 2024. <https://www.gsb.stanford.edu/insights/ab-testing-gets-upgrade-digital-age> (accessed Apr. 22, 2025).
2. "Top 64 Latest A/B Testing Statistics, Data, and Trends," *9cv9 Career Blog*, Feb. 26, 2025. <https://blog.9cv9.com/top-64-latest-a-b-testing-statistics-data-and-trends/> (accessed Apr. 23, 2025).
3. P. Guha, "30 Key A/B Testing Statistics: A Comprehensive Guide," *VWO*, Oct. 25, 2024. <https://vwo.com/blog/ab-testing-statistics/> (accessed Apr. 24, 2025).
4. "Why 94% of Financial Spreadsheets Contain Errors," *NextProcess*, Apr. 04, 2025. <https://www.nextprocess.com/uncategorized/why-94-of-financial-spreadsheets-contain-errors-and-what-it-costs-you/> (accessed Apr. 25, 2025).
5. P. Barnhurst, "How an Excel error cost JP Morgan \$6 billion," *LinkedIn*, Jul. 19, 2024. https://www.linkedin.com/posts/thefpandaguy_how-an-excel-error-cost-jp-morgan-6-billion-activity-7220200222088470528-fd8u (accessed Apr. 26, 2025).
6. "How a Tech Innovator Reduced Test Cycle Time," *Provar*, Apr. 10, 2025. <https://provar.com/case-study/how-this-technology-company-reduced-test-cycle-time/> (accessed Apr. 26, 2025).
7. R. Kohavi and R. Longbotham, "Online Controlled Experiments and A/B Testing," *Encyclopedia of Machine Learning and Data Mining*, pp. 1–8, Jan. 2016, doi: https://doi.org/10.1007/978-1-4899-7502-7_891-1.
8. "Study finds 94% of business spreadsheets have critical errors," *Higher Education Press*, Aug. 13, 2024. <https://phys.org/news/2024-08-business-spreadsheets-critical-errors.html> (accessed Apr. 26, 2025).
9. B. Zobel, "The \$6 Billion Excel Error," *Prosper Spark*, Oct. 2024. <https://www.prosperspark.com/the-6-billion-excel-error/> (accessed Apr. 27, 2025).
10. "Kafka use cases: Real-world applications," *Statsig*, Aug. 13, 2024. <https://www.statsig.com/perspectives/kafka-use-cases-applications> (accessed Apr. 28, 2025).
11. K. Smith, "Introducing the 2022 State of Feature Management," *Launch Darkly*, 2022. <https://launchdarkly.com/blog/state-of-feature-management-2022/> (accessed Apr. 28, 2025).
12. "The data-driven enterprise of 2025," McKinsey, 2022. Accessed: Apr. 30, 2025. [Online]. Available: <https://www.mckinsey.com/~media/mckinsey/business%20functions/mckinsey%20analytics/our%20insights/the%20data%20driven%20enterprise%20of%202025/the-data-driven-enterprise-of-2025-final.pdf>
13. "Knowledge Base," *Geteppo*, 2025. <https://docs.geteppo.com/experiment-analysis/reporting/knowledge-base/> (accessed May 01, 2025).
14. "The state of feature management," *Launch Darkly*, 2022. <https://launchdarkly.com/state-of-feature-management/> (accessed May 02, 2025).
15. J. Koshy, "Kafka Ecosystem at LinkedIn," *LinkedIn*. <https://www.linkedin.com/blog/engineering/open-source/kafka-ecosystem-at-linkedin> (accessed May 03, 2025).
16. A. Sellers, "Key Takeaways from Confluent's 2024 Data Streaming Report," *Confluent*, Jun. 05, 2024. <https://www.confluent.io/blog/2024-data-streaming-report/> (accessed May 04, 2025).
17. T. Chan, "You don't need large sample sizes to run A/B tests," *Statsig.com*, Dec. 24, 2024. <https://www.statsig.com/blog/you-dont-need-large-sample-sizes-ab-tests> (accessed May 05, 2025).
18. "Statistical significance," *Optimizely*, Feb. 05, 2025. <https://support.optimizely.com/hc/en-us/articles/4410284003341-Statistical-significance> (accessed May 08, 2025).
19. "Rollout and rollback features," *Optimizely*.

- <https://docs.developers.optimizely.com/full-stack-experimentation/v2.0/docs/rollout-and-rollback-features> (accessed May 09, 2025).
20. L. Pekelis, "The story behind our Stats Engine," *Optimizely*, Jan. 20, 2015.
<https://www.optimizely.com/insights/blog/statistics-for-the-internet-age-the-story-behind-optimizelys-new-stats-engine/> (accessed May 10, 2025).
 21. "A beginner's guide to Bayesian experimentation," *Statsig*, Sep. 16, 2024.
<https://www.statsig.com/perspectives/beginner-guide-bayesian-experimentation> (accessed May 10, 2025).
 22. F. Siddiqi, "ML Platform Meetup: Infra for Contextual Bandits and Reinforcement Learning," *Medium*, Oct. 18, 2019.
<https://netflixtechblog.com/ml-platform-meetup-infra-for-contextual-bandits-and-reinforcement-learning-4a90305948ef> (accessed May 11, 2025).
 23. A. Jetli, "Making Uber's Experiment Evaluation Engine 100x Faster," *Uber Blog*, Oct. 03, 2024.
<https://www.uber.com/en-GB/blog/making-ubers-experiment-evaluation-engine-100x-faster/> (accessed May 12, 2025).
 24. R. Chang, "How Airbnb Achieved Metric Consistency at Scale," *The Airbnb Tech Blog*, Aug. 05, 2021.
<https://medium.com/airbnb-engineering/how-airbnb-achieved-metric-consistency-at-scale-f23cc53dea70> (accessed May 13, 2025).
 25. "Use minimum detectable effect when designing an experiment," *Optimizely*, 2024.
<https://support.optimizely.com/hc/en-us/articles/4410288881293-Use-minimum-detectable-effect-when-designing-an-experiment> (accessed May 14, 2025).
 26. M. Stewart, "Sequential Testing on Statsig," *Statsig*.
<https://www.statsig.com/blog/sequential-testing-on-statsig> (accessed May 15, 2025).
 27. "Creating Experiment Analyses," *Geteppo*.
<https://docs.geteppo.com/experiment-analysis/configuration/> (accessed May 17, 2025).
 28. N. Donovan, "The role of experimentation at Booking.com," *Booking*, Sep. 02, 2019.
<https://partner.booking.com/en-us/click-magazine/industry-perspectives/role-experimentation-bookingcom> (accessed May 16, 2025).
 29. "7 Reasons Percentage Rollouts Reduce Deployment Risk," *Launch Darkly*, Nov. 15, 2022.
<https://launchdarkly.com/blog/how-percentage-rollouts-minimize-deployment-risks/> (accessed May 19, 2025).
 30. M. Watson, "The Ultimate Guide to Feature Flags: Implementation Strategies for Enterprise Applications," *Full Scale*, Jan. 27, 2025.
<https://fullscale.io/blog/feature-flags-implementation-guide/> (accessed May 19, 2025).