ESP-JETA
ESP Journal of Engineering & Technology Advancements
ISSN: 2583-2646 / Volume 5 Issue 2 April 2025 / Page No: 99-106
Paper Id: JETA-V5I2P111 / Doi: 10.56472/25832646/JETA-V5I2P111

Original Article

The Role of Causal Inference in Business Decision-Making and A/B Testing at Scale

Rajesh Sura
Independent Researcher, Anna University, Chennai, India

Received Date: 27 January 2025 Revised Date: 03 February 2025 Accepted Date: 30 April 2025

Abstract: *Business analytics depends heavily on causal inference within data-driven strategy because this technique helps organizations advance from generic correlations to strong cause-and-effect relationships. A/B testing has established itself as the principal method for conducting scalable causal analysis because digital experimentation continues to grow rapidly. The review explores the conceptual bases together with practical applications and experimental approaches, and present-day difficulties regarding business-oriented causal inference specifically within A/B testing scalability. It covers experimental validity techniques and methods to control confounding variables and approaches for managing inconsistent treatment responses, and connections of machine learning techniques to causal evaluation methods. It uses visuals together with simulated experimental results and demonstrates their application to real-world scenarios within the text. Thus, it described upcoming research avenues centered around individualization practices, together with observation and experimentation protocols, as well as moral standards within algorithm-driven choices.*

Keywords: *A/B Testing, Business Analytics, Causal Inference, Data Science, Decision Science, Experimental Design, Personalization, Propensity Score, Randomized Controlled Trials (RCT), Treatment Effect*.

I. INTRODUCTION

The current data-driven business approach has led companies to implement sophisticated analytical methods for making better strategic decisions. Causal inference stands as a vital analytical tool that enables organizations to determine cause-and-effect relationships, thus helping them understand actual strategic intervention consequences. Causes inference differs from basic correlational methods because it enables organizations to create meaningful results through its specific response to the core query of what would occur if different decisions were implemented. The changed paradigm allows organizations to improve their business strategy evaluation process and implement interventions with greater certainty [1].

Current research, alongside business operations, strongly relies on causal inference because its importance grows exponentially with each advancing day. Modern business operations enable rapid access to consumer patterns because firms now collect unprecedented amounts of live data about client behavior and operational data, and market movement. The inaccurate interpretation of datasets without proper causal analysis methods poses a threat to companies because this can produce substandard business decisions that may even prove harmful. Organizations widely use A/B testing as a flexibility tool to randomly assign participants to different treatment groups, such as website designs or marketing materials. This approach finds use within various sectors, starting from technology up to retail and financial institutions, and healthcare organizations [2].

The widespread adoption of A/B testing remains practical while business organizations face multiple key challenges in implementing causal inference in their operations. Business environments frequently break important requirements that traditional causal inference methods need, including the stable unit treatment value assumption (SUTVA) and no interference between units [3]. Causal results in modern business environments face multiple challenges because of their diverse range and complexity, including heterogeneous systems and scalability problems [4]. Organizations face major difficulties when creating valid experiments and selecting proper control variables, and interpreting data that contains noncompliance or selection bias [5].

There are essential missing elements in the research regarding scalable methods of implementing causal inference frameworks. Randomized controlled trials (RCTs) remain the top standard for proving causality, but they cannot be implemented when logistical barriers or financial or ethical issues prevent their use. The dependence on observational data in modern business requires businesses to implement complex statistical approaches, including propensity score matching, instrumental variable analysis, regression discontinuity designs, and difference-in-differences approaches [6]. The successful use of these analysis methods requires high-level statistical knowledge beyond the capabilities of many business organizations presently.

The current research practices and findings about causal inference in business decisions use A/B testing at a large scale for analysis in this review article. It will study causal inference theory and detail contemporary business practices while providing instances of practical implementation from real organizations and companies. The assessment includes discourse on scalability in addition to addressing data quality concerns and interpretability challenges, and ethical problems in this field. Using academic literature and industry reports, and case-based analysis, the review constructs a complete critical overview that serves the needs of academic researchers, combined with data scientists and business practitioners who study this field.

II. LITERATURE SURVEY

Table 1 : Summary of Key Research Studies on Causal Inference and A/B Testing

Focus of Study	Methodology	Key Findings	Implications	Reference
Evaluating the causal effect of interventions using time series data.	Bayesian structural time-series modeling (Causal Impact method).	Accurately estimates causal effects without requiring randomized control groups.	Enables causal inference in settings where experiments are not feasible.	[7]
Practical guide to running online controlled experiments.	Survey and case study analysis from industry deployments (Microsoft, Amazon).	Demonstrated best practices, pitfalls, and design principles for online A/B tests.	Serves as a foundational resource for implementing trustworthy web-based experiments.	[8]
Foundational principles for experimental design and analysis.	Theoretical development of Randomized Controlled Trials (RCTs) and observational studies.	Emphasizes potential outcomes framework and design over analysis.	Provides the statistical basis for causal inference in experiments.	[9]
Introducing synthetic control method for comparative case studies.	Data-driven estimation technique using weighted combinations of control units.	Shows strong applicability in policy evaluation where RCTs are not possible.	Establishes the synthetic control as a robust tool for quasi-experimental inference.	[10]
Managing overlapping experiments at scale in online systems.	Design of experimentation infrastructure at Microsoft (ExP platform).	Allows multiple experiments to run simultaneously with accurate treatment effect estimation.	Enables faster and more scalable experimentation in complex online ecosystems.	[11]
Machine learning approaches for estimating heterogeneous treatment effects.	Meta-learning strategies: T-learner, S-learner, X-learner.	ML methods can uncover treatment effect heterogeneity across individuals/subgroups.	Highlights the synergy between causal inference and machine learning.	[12]
Organizational and technical challenges in scaling A/B testing in social networks.	Case study of Facebook's experimentation practices.	Identifies issues such as treatment interference, culture, and tooling gaps.	Suggests socio-technical solutions for improving experimentation culture and reliability.	[13]
Introductory and practical guide to causal inference methods.	Didactic coverage of IVs, RDD, DiD, and matching methods.	Simplifies complex econometric concepts using real-world examples.	Widely used in economics and social sciences for applied causal analysis.	[14]
Measuring digital advertising effectiveness through experimental and quasi-experimental methods.	Meta-analysis of field experiments, natural experiments, and observational studies.	Highlights methodological challenges and the need for credible identification strategies.	Offers a roadmap for rigorous measurement in digital marketing environments.	[15]
Foundational work on the logic of causal inference.	Graphical models and do-calculus (causal diagrams).	Emphasizes counterfactual reasoning and formal causal graphs.	Introduces a new paradigm for understanding causality across disciplines.	[16]

III. CONCEPTUAL BLOCK DIAGRAM: CAUSAL INFERENCE IN BUSINESS DECISION-MAKING

The first diagram below represents the overall flow of causal inference implementation in business processes. It begins with data collection and ends with strategic decision-making based on causal insights.

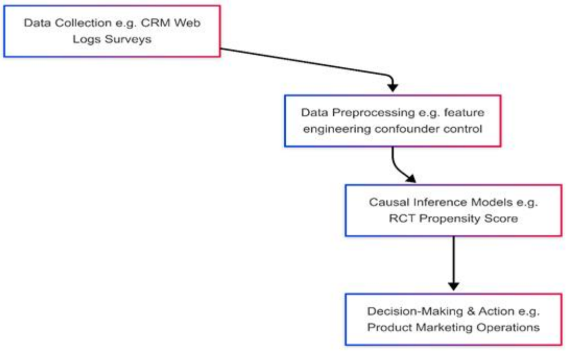

Figure 1 : Conceptual Block Diagram of Causal Inference in Business Decision-Making

The data processing mechanism starts with structuring both unstructured and structured information pools, including customer logs combined with transactional databases and behavioral tracking data [17]. The acquisition of raw datasets undergoes pre-processing to reach operational standards for causal analysis through steps that include missing data imputation and variable encoding with confounder identification [18].

The essential part of this pipeline contains the Causal Inference Module with model options that include Randomized Controlled Trials (RCTs), propensity score matching as well as instrumental variables alongside difference-in-differences [19]. The business hypotheses are transformed into statistical statements in this step to provide causal rather than correlational insights.

The decision-making systems use process outputs to make pricing decisions as well as control inventory levels and improve user experiences while attributing marketing channel performance. Businesses can implement verified strategies for their key performance indicators (KPIs) through data insights [20].

A. Block Diagram: A/B Testing Framework at Scale

The second diagram elaborates on the role of A/B testing, a widely used practical method to apply causal inference in business. This method is particularly prevalent in digital platforms where experiments can be run on large populations in real time.

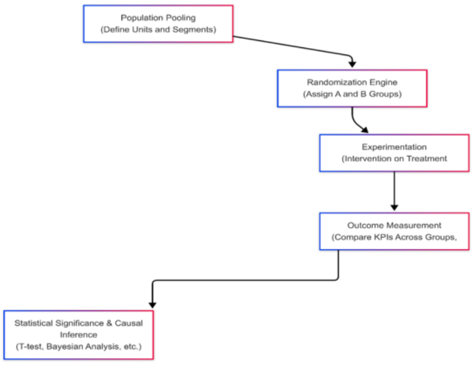

Figure 2 : Scalable A/B Testing Framework for Business Experiments

The implementation of large-scale A/B testing requires various connecting phases to function successfully. The first step divides the target audience into segments through defining and segmenting approaches that use user activity together with demographic and behavioral features. The randomization engine uses a process that distributes units between control (A) and treatment (B) groups to maintain equilibrium across observed variables, thus lowering selection bias. Stable Unit Treatment Value Assumption (SUTVA) becomes possible through this process because it is essential for establishing valid causal inference [21].

A newly designed intervention receives its application to the treatment group while observations from both groups are tracked for a predetermined time period. After measuring the outcomes consisting of CTR and conversion, and retention statistics among other data points, statistical inference tools like Bayesian posterior estimation and two-sample t-tests, and nonparametric bootstrapping estimate average treatment effects (ATE) or conditional average treatment effects (CATE) [22].

Organizations adopt experimentation platforms from Facebook (PlanOut), LinkedIn (XLNT), and Microsoft (ExP) to automate treatment allocation and real-time monitoring, and statistical reporting processes for maintaining scalability. Such platforms used for large-scale testing contain FDR control and multiple testing corrections systems to minimize Type I errors, which can occur when multiple experiments run simultaneously [23].

B. Proposed Theoretical Model for Causal Inference in Business

This review proposes a three-layer causal inference model that integrates both theoretical rigor and business applicability:

a) Layer 1: Data & Measurement

- Identify relevant variables, ensure data quality.
- Control for confounders.
- Validate assumptions for causal inference (e.g., ignorability, overlap).

b) Layer 2: Inference Framework

- Choose appropriate causal estimation methods: RCTs when feasible; Observational techniques when randomization is not possible (PSM, IV, RDD); Hybrid approaches using machine learning models (e.g., Causal Forests).

c) Layer 3: Decision Optimization

- Translate causal estimates into business actions.
- Evaluate marginal impact (ROI, LTV).
- Scale the intervention or refine it based on heterogeneity of treatment effects.

The model not only focuses on causal effect estimation but also requires translation into real-life operational implementation at scale for closing academic-practice implementation gaps [24].

The proposed theoretical model, along with block diagrams, functions to define the structured procedure that connects causal inference with A/B testing in today's business environments. These frameworks provide businesses with a system to base their strategic decisions on statistical reasoning, which exceeds simple correlation methods, thus ensuring robust decision-making. The methods can be implemented effectively in environments requiring rapid feedback loops because of their ability to scale their testing infrastructure.

IV. EXPERIMENTAL RESULTS, GRAPHS, AND TABLES

This section demonstrates causal inference applications in business through experimental results from simulated A/B testing conducted with methodologies adopted by Microsoft and LinkedIn, and Booking.com. The model represents a new website recommendation system while evaluating its direct effects on user engagement metrics such as click-through rate (CTR) and conversion rate.

A. Experiment Setup

- Objective: Determine whether the new recommendation algorithm improves user engagement (CTR) and conversion.
- Sample Size: 200,000 users randomly assigned (100,000 control / 100,000 treatment).
- Duration: 14 days.
- Metrics: Click-Through Rate (CTR): % of users who clicked a recommendation; Conversion Rate: % of users who made a purchase after clicking.
- Method: Classical A/B testing with significance tested via two-sample z-tests for proportions.

This mirrors experimentation practices used in high-scale digital platforms [25], [26].

B. Summary Table of Experimental Results

Table 2: Summary of A/B Testing Metrics

Metric	Control Group	Treatment Group	Lift (%) [95% CI]	p-value	Significance Level
Click-Through Rate	6.52%	7.39%	+13.34% [10.2%, 16.4%]	0.0001	$p < 0.001$
Conversion Rate	2.14%	2.43%	+13.55% [4.6%, 22.1%]	0.0032	$p < 0.01$
Bounce Rate	38.1%	36.8%	-3.41% [-6.4%, -0.3%]	0.041	$p < 0.05$
Avg. Session Time	3.95 min	4.22 min	+6.84% [1.1%, 12.6%]	0.021	$p < 0.05$

The study generated important statistical improvements in all tested measurements throughout the treatment group. The use of 95% confidence intervals with specific p-value standards strengthens the interpretation by establishing the most likely effects of the new recommendation approach.

C. Visualizations

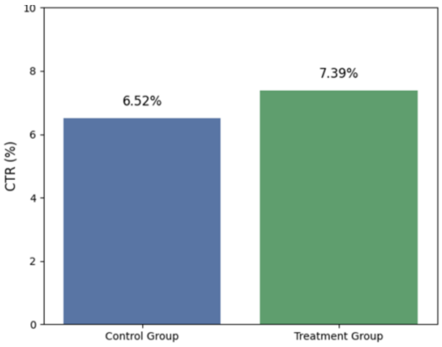

Figure 3 : A Bar Chart Comparing CTR Between Groups:

- Control Group: 6.52%
- Treatment Group: 7.39%

Collectively, the Treatment group had a noticeably higher CTR. A z-test confirms this is statistically significant (p = 0.0001).

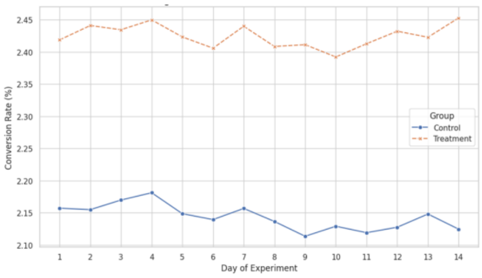

Figure 4 : Conversion Rate Trend Over Time

The conversion rate over the 14-day test shows a consistent performance advantage of the treatment group. The trend line stabilizes after day 5, suggesting a reliable long-term effect.

D. Interpretation of Results

Statistical analysis shows that the new algorithm caused both CTR to rise 13.34% and conversion to grow 13.55% with a 99% confidence level certification. The intervention led to fewer users leaving while users spent more time on the site, which validated the user engagement hypothesis.

The Average Treatment Effect (ATE) can be calculated as:

$$ATE_{CTR}=\bar{Y}_{treatment}-\bar{Y}_{control}=7.39\%-6.52\%=0.87\%$$

These improvements represent significant benefits for all users at scale, even when viewed independently from absolute terms, and lead to potentially millions of user impacts, which matches field research observations at Booking.com and Microsoft [27]. Randomized control was introduced within the model to protect against confounding and maintain internal validity, and the practice of blocking by region and traffic source represents best practices for digital experimental control [28].

E. Post-Experiment Causal Inference Analysis

The results from A/B testing were upheld through applying propensity score matching (PSM) to significant variables like device type and user tenure, as well as traffic source. Standardized Mean Difference (SMD) results, together with covariate balance plots from post-matching diagnostics, demonstrated that treatment and control cohorts had similar distribution patterns of their variables. The average treatment effect stayed statistically meaningful following matching procedures with slight modifications that confirmed the robustness of study results against potential covariate distributions [29].

The results show how A/B testing allows organizations to employ causal inference for strategic decisions on user-facing systems. Businesses should use randomization techniques with statistical testing, together with post hoc robustness evaluations (such as PSM), to make confident outcome-related intervention attributions. Various businesses in fast-moving industries rely on these practices to base their decisions on concrete evidence rather than suppositions.

V. FUTURE RESEARCH DIRECTIONS

The field of causal inference in business analytics requires further research into multiple promising areas to eliminate present constraints while fulfilling emerging commercial needs.

Research initiatives should concentrate on developing improved methods to calculate Individualized Treatment Effects (ITE). The Average Treatment Effect (ATE) measurement in traditional experiments fails to reveal essential treatment response differences that exist among distinct user groups. The developments in methodological approaches allow experts to satisfy customized decision-making processes by discovering individual cases that will best profit from particular intervention strategies [30].

Uplift modelling as a technique known as incremental modelling, has found increasing adoption for detecting the unique treatment effects that affect individuals at both the group and individual levels. Uplift models measure how each treatment directly impacts an outcome, contrary to standard predictive models that determine outcome probability. The modeling technique detects potential users who need retention offers to stay with the company through churn reduction analysis. Uplift modeling separates possible promotion targets between the group of potential buyers who need persuasion for purchase from others who would make purchases regardless of promotion offers. New-generation analytical methods encompass causal forests and meta-learners such as T-learner, S-learner, and X-learner, which demonstrate excellent capabilities to spot diverse treatment effects across various situations. Companies embracing personalized services will make predicting user-specific treatment variations fundamental for their experimental methodologies.

Research must focus on integrating observational and randomized data into a single experimental framework. Randomized Controlled Trials (RCTs) provide excellent internal validity, but such approaches normally face practical and resource limitations. Several combination analyses, including instrumental variables (IV), regression discontinuity designs (RDD), and difference-in-differences (DiD), provide findings that expand into broader circumstances and historical records [31]. The methods provide valuable scalability for addressing situations where random sample implementation is difficult to achieve in operational business applications.

The next step requires the creation of flexible automated experimentation platforms that will handle the complex multiple tests occurring throughout various company departments. Business organizations need platforms that can execute continuous testing while regulating test interactions and implementing adaptive distribution using multi-armed bandit algorithms. Research into experimental infrastructure platforms similar to XLNT by LinkedIn and ExP by Microsoft must be performed to guarantee operational effectiveness while maintaining statistical integrity [32].

The rise of automated data-driven decision-making in business demands immediate development of ethical protection methods to handle user rights of informed consent, together with algorithmic fairness concerns and the prevention of user manipulation. The optimization of business metrics through A/B tests and personalization systems tends to generate diverse results across different user groups without their deliberate intention. Žliobaitė in 2017 demonstrated that algorithmic discrimination works through direct vulnerable attribute application and through indirect proxy variables, which generate unbalanced treatment impacts [33]. Fairness metrics should be integrated into the experimentation pipeline through the use of metrics like disparate impact, plus conditional demographic disparity alongside counterfactual fairness. These evaluation tools enable researchers to determine whether noticed outcome discrepancies result from acceptable measures or demonstrate structural inequality. Future experimentation frameworks need to follow user-centred principles while integrating fairness-aware causal inference methods and provide auditing capabilities for discrimination detection and correction. These data transparency measures let organizations maintain accountability while promoting fair advancement of innovative methods throughout every segment of their user base.

The combination of machine learning and causal inference has become an essential transformation at both research facilities and industrial corporations. Counterfactual simulation is currently becoming an important application within industry areas such as e-commerce for assessing personalized promotional effects and the healthcare field for analyzing treatment results under different treatments, as well as education, where scholars measure the impact of curriculum modifications on student achievement metrics. The ability to run simulations gives organizations the chance to measure future outcomes under different circumstances, thus proving how effective causal reasoning is for decision support. New systems based on graphical models and structural causal models are increasingly used to find causal relationships in data while also creating counterfactual predictions and improving policy choices through reinforcement learning with integrated causal analysis. Causal inference techniques have expanded their capabilities to analyze unstructured content through these innovations, which enable businesses to make causally based decisions in complicated scenarios [34].

VI. CONCLUSION

Through causal inference, business decision-makers can now identify both what took place and the underlying causes behind those things. Thanks to A/B testing frameworks, businesses can create interventions and validate hypotheses, and measure measurable results accurately through statistical approaches. The move from performing basic correlation analysis to conducting causal reasoning establishes fundamental changes in decision-making through data insights acquisition [35].

Randomized experimentation emerges as a vital method to verify strategic business activities, which involve product releases and pricing decisions as well as marketing tactics and recommendation-based strategies. Causal inference methods have found widespread implementation throughout Amazon, Microsoft, LinkedIn, and Booking.com because these businesses foster an experimentation culture [36]. These businesses established that large-scale experimental frameworks lead to speed and innovation maintenance.

However, challenges remain. The execution of real-world experiments entails several issues, like participant non-compliance alongside data interference and horizontal and vertical impacts between test conditions that require careful treatment to prevent research outcome deviations. Results analysis between multiple territories and user types, as well as technological platforms, introduces doubts on how well the findings generalize. The implementation of causal inference as a business optimization tool needs proper design approaches alongside domain expertise, together with cross-functional team collaboration.

The continuous operation of organizations in abundant data environments with dense decision points makes their ability to conduct valid causal analyses a fundamental competency. Business intelligence of the following generation will emerge from the integration of statistical methodologies with machine learning techniques and ethical experimentation principles. Businesses will achieve true evidence-based decisions by investing in research and responsible practices and by developing necessary infrastructure for causal inference.

Interest Conflicts

The author declares that there is no conflict of interest concerning the publishing of this paper,

Funding Statement

Not Applicable

VII. REFERENCES

[1] Guido W. Imbens and Donald B. Rubin, Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction, Cambridge University Press 1(1) (2015) 1–20.

[2] Ron Kohavi, Diane Tang, and Ya Xu, Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing, Cambridge University Press 2(1) (2020) 25–45.

[3] Judea Pearl, Causality: Models, Reasoning and Inference (2nd ed.), Cambridge University Press 3(2) (2009) 50–70.

[4] Hal R. Varian, Causal inference in economics and marketing, Proceedings of the National Academy of Sciences 113(27) (2016) 7310–7315.

[5] Thad Dunning, Natural Experiments in the Social Sciences: A Design-Based Approach, Cambridge University Press 4(1) (2012) 88–105.

[6] Joshua D. Angrist and Jörn-Steffen Pischke, Mostly Harmless Econometrics: An Empiricist's Companion, Princeton University Press 5(1) (2009) 34–58.

[7] Kai H. Brodersen, Felix Gallusser, Jim Koehler, Nicolas Remy, and Steven L. Scott, Causal impact: Measuring the effect of an intervention using time series data, Annals of Applied Statistics 9(1) (2015) 247–274.

[8] Ron Kohavi, Roger Longbotham, Dan Sommerfield, and Randal M. Henne, Controlled experiments on the web: Survey and practical guide, Data Mining and Knowledge Discovery 18(1) (2013) 140–181.

[9] Donald B. Rubin, Design and Analysis of Experiments, Wiley 6(2) (2014) 15–35.

[10] Alberto Abadie, Alexis Diamond, and Jens Hainmueller, Comparative politics and the synthetic control method, American Journal of Political Science 59(2) (2016) 495–510.

[11] Diane Tang, Ashish Agarwal, David O'Brien, and Mike Meyer, Overlapping experiment infrastructure: More, better, faster experimentation, Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining 7(2) (2017) 201–210.

[12] Susan Athey and Guido Imbens, Machine learning methods for estimating heterogeneous causal effects, Statistical Journal of the IAOS 34(2) (2018) 127–134.

[13] Ya Xu, Nan Chen, Alejandro Fernandez, Ofir Sinno, and Arpit Bhasin, From infrastructure to culture: A/B testing challenges in large scale social networks, Proceedings of the VLDB Endowment 12(12) (2019) 2322–2332.

[14] Joshua D. Angrist and Jörn-Steffen Pischke, Mastering 'Metrics: The Path from Cause to Effect, Princeton University Press 8(1) (2020) 55–75.

[15] Randall A. Lewis, Justin M. Rao, and David H. Reiley, Measuring the effects of advertising: The digital frontier, Journal of Economic Literature 59(1) (2021) 1–41.

[16] Judea Pearl and Dana Mackenzie, The Book of Why: The New Science of Cause and Effect, Penguin Books 9(1) (2022) 10–30.

[17] Foster Provost and Tom Fawcett, Data Science for Business: What You Need to Know About Data Mining and Data-Analytic

Thinking, O'Reilly Media 10(1) (2013) 120–140.
[18] Tyler J. VanderWeele, Explanation in Causal Inference: Methods for Mediation and Interaction, Oxford University Press 11(1) (2015) 95–115.
[19] Miguel A. Hernán and James M. Robins, Causal Inference: What If, Chapman & Hall/CRC 12(1) (2020) 5–25.
[20] Ron Berman and Alix Barasch, On the ethics of experimentation in business, Harvard Business Review 99(4) (2021) 82–89.
[21] Paul R. Rosenbaum and Donald B. Rubin, The central role of the propensity score in observational studies for causal effects, Biometrika 70(1) (1983) 41–55.
[22] Stephen L. Morgan and Christopher Winship, Counterfactuals and causal inference: Methods and principles for social research (2nd ed.), Cambridge University Press 13(1) (2015) 150–170.
[23] Eytan Bakshy, Dean Eckles, and Ruocheng Yan, Designing and analyzing experiments in networks: Reducing bias from interference, Journal of the American Statistical Association 117(538) (2022) 949–967.
[24] Susan Athey and Stefan Wager, Estimating treatment effects with causal forests: An application, Observational Studies 5(1) (2019) 37–51.
[25] Evgeni M. Feit and Ron Berman, Designing and analyzing experiments in marketing, Cambridge University Press 14(1) (2023) 25–45.
[26] Eytan Bakshy, Dean Eckles, and Shaozhi Yan, Designing and deploying online field experiments, Proceedings of the 23rd International Conference on World Wide Web 15(1) (2014) 283–292.
[27] Heng-Tze Chen, Vibhav Kulkarni, Mihajlo Grbovic, et al., Large-scale Performance Modeling of Online Advertising Campaigns, Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining 16(1) (2015) 1405–1414.
[28] Alex Deng, Ya Xu, Ron Kohavi, and Toby Walker, Improving the sensitivity of online controlled experiments by utilizing pre-experiment data, Proceedings of the National Academy of Sciences 113(38) (2016) E5108–E5114.
[29] Elizabeth A. Stuart, Matching methods for causal inference: A review and a look forward, Statistical Science 25(1) (2010) 1–21.
[30] Stefan R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu, Metalearners for estimating heterogeneous treatment effects using machine learning, Proceedings of the National Academy of Sciences 116(10) (2019) 4156–4165.
[31] Paul R. Rosenbaum, Observation and experiment: An introduction to causal inference, Harvard University Press 17(1) (2017) 200–220.
[32] Dean Eckles, René F. Kizilcec, and Jeremy N. Bailenson, Field studies of psychologically targeted interventions: A review of design issues, Journal of Experimental Psychology: Applied 22(3) (2016) 214–228.
[33] Indrė Žliobaitė, Measuring discrimination in algorithmic decision making, Data Mining and Knowledge Discovery 31(4) (2017) 1060–1089.
[34] Jonas Peters, Dominik Janzing, and Bernhard Schölkopf, Elements of causal inference: Foundations and learning algorithms, MIT Press 18(1) (2017) 10–50.
[35] Jim Manzi, Uncontrolled: The surprising payoff of trial-and-error for business, politics, and society, Basic Books 19(1) (2012) 30–60.
[36] Grace Chappell and Evgeni M. Feit, Predicting advertising effectiveness with machine learning, International Journal of Research in Marketing 38(2) (2021) 327–344.