FERI INFORMATION

Special Issue on Frontiers of Causal Inference

– Explanatory Notes (Part I) on SBI Research Review Vol.8 –

Yutaka Soejima, Principal Research Fellow of SBI Financial and Economic Research Institute

Frontiers of Causal Inference

The eighth volume of the SBI Research Review (available only in Japanese) focuses on the theme "The Frontiers of Causal Inference." Causal inference refers to a framework and a set of analytical methods aimed to clarify cause-and-effect relationships, that is, to identify what has driven a particular outcome or what consequences may result from a given action or policy.

For example, when public authorities implement vaccination programs, provide direct income support to households, cut taxes, raise the minimum wage, increase individuals' out-of-pocket medical expenses, revise labor contract laws for non-regular employment, promote equity investment, or impose higher tariffs, it is often difficult to determine whether the observed outcomes are truly attributable to these policy interventions or merely reflect the influence of other external factors. Likewise, in the business sphere, when companies lower consumer product prices, reduce product size or quantity (a practice commonly referred to as shrinkflation), distribute coupons, launch targeted advertising, redesign store layouts, conduct AI training programs for employees, add new app features, expand the number of outside directors, or carry out IPOs, MBOs, or M&A transactions, isolating and accurately measuring the impact of these actions is extremely challenging due to the presence of confounding variables.

The central challenge of causal inference lies in the impossibility of observing both the outcome with an intervention and the outcome without it for the same economic agent or entity under identical conditions. This dilemma is commonly referred to as *the fundamental problem of causal inference*. Since we cannot construct parallel universes in which all other factors are held constant and compare a world with intervention to one without intervention, researchers have devoted extensive efforts to developing methodologies that enable causal testing despite this inherent limitation

1. Historical Development and Core Methodologies

The concept of a counterfactual outcome refers to a hypothetical scenario—a plausible world that would have occurred if a particular policy or business strategy had, or had not, been

FERI INFORMATION (25.09.10) - Posted September 1, 2025 in Japanese

implemented. This notion is familiar to economists, and various empirical modeling techniques have been developed to estimate such counterfactuals. Consider, for example: What if the consumption tax had not been raised? What if quantitative easing had not been introduced? What if tightening measures had been adopted earlier during an asset bubble? Or what if public funds had been deployed sooner to address the non-performing loan crisis? The first two questions represent counterfactuals in which policies are assumed not to have been implemented, while the latter two assume they were. In addition, instrumental variable (IV) methods, widely employed in econometrics, have also been applied to causal inference. However, structural model-based approaches raise concerns about whether the model appropriately reflects the real-world dynamics, while IV methods raise concerns about whether suitable instruments are available. Both issues continue to pose significant challenges for empirical analysis.

Among the most prominent techniques in causal inference is the Randomized Controlled Trial (RCT), which addresses the fundamental problem through experimentation. Developed in the 1930s and applied in medicine in the 1940s, RCTs were long considered impractical in economics due to technical, ethical, and budgetary constraints. However, since the 1990s, RCTs have been increasingly applied in social policy, education, and economics, where they have played a central role in advancing development economics. Notably, the 2019 Nobel Prize in Economics was awarded for work involving RCTs in development economics. By the 2010s, RCTs came to be widely applied in labor economics, education economics, behavioral economics, and public economics.

Beyond RCTs and IV methods, a variety of other techniques have been developed in causal inference, including propensity score matching, difference-in-differences, regression discontinuity design, and inverse probability weighting. These approaches—often grouped under the umbrella of potential outcome models—facilitate aggregate-level estimation by leveraging statistical assumptions and design refinements to address the unobservability of counterfactuals. In parallel, structural causal models have also evolved. This framework expresses causal directionality between causes and effects through graphical representations, which are then translated into structural equations, thereby enabling empirically testable counterfactual analysis at the individual level. Although the modeling methodology differs, the estimation of structural causal models closely resembles the structural modeling techniques used in econometrics discussed earlier.

2. Applications in Policy and Business

The advances in causal inference outlined above have become increasingly important in the context of Evidence-Based Policy Making (EBPM) and data-driven marketing. For instance, a landmark study conducted in the United States in the early 1990s challenged the conventional view of a trade-off between raising the minimum wage and higher unemployment. Another study, using bank account data, revealed that a substantial portion of COVID-19 relief payments was not directed toward consumption. Both relied on causal inference methodologies. Causal inference has likewise been critical in assessing the impacts of a wide range of policies: whether environmental regulations such as emissions trading effectively alter corporate behavior and reduce air pollution; whether smaller class sizes contribute to improve academic performance; whether health screenings and guidance help lower medical expenditures; and whether increased police presence or the installation of surveillance cameras suppress crime rates. Across these diverse domains, EBPM continues to evolve, with causal inference serving as a central tool for empirical validation.

The business sector also has witnessed a wide array of practical applications. Companies are actively experimenting with whether recommendation algorithms on video-streaming platforms extend viewing time or reduce subscription cancellations; whether coupon distribution increases taxi ridership and revenue, or conversely, whether excessive distribution erodes profitability; whether financial product campaigns or app UI/UX enhancements stimulate purchases and revenue; and whether CRM tools, sales strategies, or training programs improve conversion rates, average customer spending, or lifetime value. In short, the sophistication of data-driven strategies has long been directly tied to competitive advantage. The extraordinary growth in the market capitalization of GAFAM companies exemplifies this trend, underscoring the dominance of data-driven business models.

At the same time, however, many policy implementations and strategic decisions continue to be guided by heuristics or managerial intuition. Given the current level of data science adoption and the persistent limitations of IT infrastructure and data readiness, it is entirely plausible that decisions in both policymaking and business execution are still being shaped by flawed or incomplete analyses.

The promotion and adoption of causal inference methods have likewise become increasingly prominent. As interest in the field grows, a wide range of specialized, introductory, and practical texts have been published. Since around 2020, numerous high-quality volumes—both translations and original works by Japanese scholars—have been released.

3. Focus of This Special Volume: Emerging Analytical Methods and the Extension to Language as a Target of Analysis

The aim of this volume is to highlight recent developments in causal inference. One of the key driving forces behind its evolution in recent years has been the rapid advancement of natural language processing (NLP), particularly the explosive progress of Large Language Models (LLMs). These technologies have enabled the use of textual data—written language—as a new source for causal analysis. (For further discussion of LLMs, see Soejima's article in Research Review vol. 6 [2024], available only in Japanese.) Beyond the statistical approaches to causal inference discussed thus far, the emergence of new methodologies and the availability of vast textual data sources represent a significant expansion of the potential of causal inference. This expansion has broadened the scope of research domains. In addition to verifying causal relationships, researchers are now examining how humans perceive causality and how such perceptions shape behavior—extending causal research into the realm of human cognition. Two papers featured in this volume embody this perspective, presenting empirical case studies that demonstrate the practical application of these methodologies.

Moreover, the growing sophistication of machine learning and AI has spurred a wave of research and business applications that integrate these technologies into causal inference. In the case of item selection in recommendation algorithms, the intervention space spans tens of thousands of songs, videos, and books—making the problem far more high-dimensional than interventions involving the adoption of a single policy. In such contexts, conducting RCTs or A/B tests is often impractical, and decisions must instead rely on observational data collected under specific policies or environmental conditions. There are also broad domains—such as healthcare—where experimental interventions are frequently infeasible. One of the papers featured in this volume addresses these decision-making challenges within the framework of off-policy evaluation, highlighting recent methodological developments.

The academic disciplines that have traditionally engaged in causal inference research span a broad spectrum. With the advent of new techniques such as AI and machine learning, and the emergence of textual data as a new source of information, the scope of academic disciplines engaged in causal inference research has expanded to encompass fields such as information processing, behavioral economics, political science, medicine, law and legal systems, economic history, comparative institutional analysis, communication science, marketing, and computational social science. The author of these explanatory notes attended the annual conference of the Japanese Society for Artificial Intelligence held in Osaka this May, where numerous sessions were dedicated to topics such as computational social science using SNS data and the application of NLP in marketing. Notably, the presenters represented a diverse array of academic disciplines and industries. In line with this diversity, the papers featured in this volume are not confined to contributions from economists. This volume includes only one paper authored by economists. The remaining contributions come from researchers with diverse disciplinary backgrounds—including engineering fields such as computer science, artificial intelligence, and machine learning, as well as education policy, statistics, NLP, and cognitive science. Each offers a distinct perspective on the study of causal inference.

4. The Potential of Causal Inference

This interdisciplinary expansion of causal inference fosters a deeper understanding of the role that causal reasoning plays in human decision-making and its broader implications for society. The internet has triggered a transformation in the production and distribution of information—a shift as profound as Gutenberg's invention of the printing press. Today, the communication and dissemination of information through social media platforms exert an influence on public opinion, behavioral change, political and economic policymaking, and the functioning of democracy that rivals, and in some cases surpasses, that of traditional media outlets or official communications from political institutions, policymakers, and corporations. The darker side of today's digital society includes the malicious manipulation of information, the rampant spread of populism, and the unauthorized appropriation and exploitation of personal privacy and security data. At the same time, there is growing anticipation for new forms of democracy enabled by digital technologies, as well as expanding possibilities for bottom-up models of social governance.

In this social environment, the importance of causal inference techniques continues to grow. To accurately identify the causes behind events occurring in society and to design effective policies and institutional frameworks, it is essential not only to carefully examine the causal relationships underlying social phenomena, but also to understand how these relationships are perceived by society and how such perceptions shape behavior. This dual perspective is a key to promoting sound social governance. The importance of this approach becomes evident even when considering specific issues such as climate change, social security, and debates over taxation.

In the business domain, causal inference holds immense potential to uncover relationships that traditional, numerically structured data often fail to reveal. Many business activities rely on unstructured textual data—such as customer feedback, service interactions, sales notes, product reviews, chat logs, social media posts, and FAQs—as a base for decision-making. Historically, such data have been employed primarily for sentiment analysis or topic extraction. Today, however, causal inference is being applied to identify how specific contexts, expressions, or messages causally shape behavior and evaluation.

The scope of applications continues to grow, encompassing the measurement of intervention effects delivered through text, audio, or video—such as advertisements or customer service interactions; the use of LLMs to process and structure large volumes of textual data to enhance the precision of causal inference; and the application of text generation to design optimized interventions.

The following section introduces the papers featured in this volume, with particular emphasis on their practical applications. For technical details, readers are encouraged to consult the individual articles (available only in Japanese) as well as the foundational research papers on which they build.

Continued in Part II: Reviews of papers in the volume.