ISSN: 2054-7404
Archives of Business Research; Vol. 14 No. 11 (2026) (151-172)
https://doi.org/10.14738/abr.1411.12658

The Dialectical Transformation of Ceteris Paribus in the Age of Big Data: Toward a Hybrid Paradigm

Avtandil Gagnidze[1] & Mikheil Tokmazishvili[2]

1. East European University, East-West University, Tbilisi, Georgia
2. Alte University, Tbilisi, Georgia

Abstract: The structural limitations of the *ceteris paribus* assumption in complex economic systems raise fundamental questions about its methodological relevance in the data-driven era. This article examines the transformation of the epistemological role of *ceteris paribus* under the conditions of Big Data, machine learning (ML), artificial intelligence (AI), and agent-based modeling (ABM). The paper argues that Big Data does not invalidate the logic of *ceteris paribus*; rather, it transforms its operationalization. The assumption shifts from theoretically imposed "constant conditions" toward empirically informed control mechanisms based on large-scale data analysis. This transformation is enabled by interpretable machine learning approaches (LIME, SHAP, PDP, ICE), nowcasting techniques, ABM, and contemporary methods of causal inference. The central thesis of the paper is that *ceteris paribus* and Big Data are not antagonistic but instead form a dialectical interaction, in which abstraction (deduction) and empirical reality (induction) are reconciled through abduction. This interaction generates an iterative cycle in which each stage contributes to the refinement of both economic theory and its empirical foundations. Within this framework, Friedman's instrumentalist methodology is reconsidered in light of the analytical possibilities offered by Big Data. The paper proposes a hybrid modeling paradigm in which economic theory provides the causal structure, while Big Data and AI/ML enable empirical validation, calibration, and refinement of this structure. To quantify this dialectical relationship, the paper introduces the Ceteris Paribus Validity Coefficient (CPV) and its extensions -CPV(BD) and AI-CPV - as instruments for evaluating the validity of *ceteris paribus* assumptions. The CPV framework provides a formal metric for assessing how well a given model approximates the "other things equal" condition, with higher values indicating stronger empirical grounding and reduced reliance on untested abstractions.

Keywords: *ceteris paribus*, Big Data, machine learning, artificial intelligence, agent-based modeling, nowcasting, dialectics, deduction–induction–abduction, hybrid paradigm, causal analysis.

INTRODUCTION

The analysis of the methodological foundations of *ceteris paribus* and the structural limitations of its application, including omitted variable bias (OVB), critiques arising from behavioral economics, the unrealistic assumption of "other conditions remaining constant" in complex economic systems, and the limitations of Friedman's instrumentalism, demonstrates that *ceteris paribus* can no longer be treated as an a priori and untested postulate. Its application requires explicit consideration of contextual conditions, empirical evidence, and the identification of underlying causal relationships.

At the same time, recent decades have witnessed a substantial expansion of economic analytical capabilities driven by Big Data analytics, machine learning (ML), and artificial intelligence (AI). Contemporary researchers can process and analyze thousands of interconnected variables simultaneously, detect complex patterns, and explore relationships that were previously difficult to identify within traditional modeling frameworks based primarily on theoretical abstraction and simplification.

Against this background, the paper addresses the central research question: **Under what conditions, and to what extent, can Big Data methods empirically validate, operationalize, or replace the *ceteris paribus* assumption?**

The central hypothesis of the study is that Big Data does not eliminate the logic of *ceteris paribus*; rather, it enables its empirical operationalization and contextual validation. This transformation becomes particularly significant within a hybrid modeling paradigm, where deduction, induction, and abduction interact through an iterative process of theoretical refinement and empirical verification.

The objective of this paper is to develop a conceptual framework for understanding the evolving relationship between *ceteris paribus* and Big Data. The scientific contribution of the study consists of three main elements: (1) reconceptualizing *ceteris paribus* not as a fixed a priori assumption but as an element of a dynamic dialectical process; (2) proposing **a spiral model** of the interaction between deduction, induction, and abduction; and (3) reconsidering Friedman's instrumentalist methodology in light of the analytical possibilities provided by Big Data.

The paper proceeds as follows. The second section examines the Big Data revolution, its five fundamental characteristics (5Vs), and the role of machine learning (ML), artificial intelligence (AI), and agent-based modeling (ABM) in contemporary economic analysis. It further discusses the methodological implications of nowcasting, causal inference approaches, dependence on historical data, along with related challenges, including Lucas's critique, the risks of overfitting, and ethical considerations. Particular attention is given to the conceptualization of the dialectical unity of deduction, induction, and abduction as a dynamic mechanism of economic knowledge formation.

Building on this theoretical framework, the paper introduces three quantitative assessment instruments - the **Ceteris Paribus Validity Coefficient (CPV)**, its Big Data extension **CPV(BD)**, and the AI-integrated version **AI-CPV -** designed to evaluate the validity, applicability, and empirical robustness of *ceteris paribus* assumptions within data-driven economic models.

THE BIG DATA REVOLUTION: A PARADIGM SHIFT IN ECONOMIC ANALYSIS

In the first decades of the twenty-first century, the methodology of economic modeling has undergone a profound structural transformation. One of the principal drivers of this transformation is the emergence of Big Data, which represents more than merely an increase in the volume of available information. Rather, it constitutes a new analytical paradigm that reshapes the processes of economic data collection, processing, and knowledge generation.

Under the conditions of the digital economy, the analysis of economic processes increasingly relies on large-scale, diverse, and high-frequency data generated through digital platforms, sensors, transactions, and other information systems. Big Data expands the analytical capabilities of economics by addressing several limitations traditionally associated with econometric modeling, including restricted sample sizes, incomplete information, and difficulties in capturing complex interactions among variables. However, its contribution lies not only in the quantity of data but also in its variety, granularity, temporal resolution, and potential for real-time analysis.

Consequently, Big Data enables the development of more dynamic, complex, and multidimensional representations of economic processes. At the same time, it does not eliminate the fundamental methodological challenges of economic analysis, such as causal identification, model specification, and omitted variable bias, but provides new tools for addressing and reassessing these challenges within a broader empirical framework.

THE MAIN CHARACTERISTICS OF BIG DATA AND THEIR RELATIONSHIP WITH *CETERIS PARIBUS*

The methodological significance of Big Data becomes more apparent when its main characteristics are examined in direct relation to the *ceteris paribus* assumption. In the literature, Big Data is traditionally described through five interconnected dimensions: **Volume, Velocity, Variety, Veracity, and Value** [1]. Each of these dimensions influences, in different ways, the applicability and operationalization of the principle of "other things remaining equal."

Volume. The substantial increase in the number of observations reduces the need for excessive abstraction from potentially relevant variables that previously could not be incorporated due to limited sample sizes. By expanding the empirical basis of analysis, Big Data allows researchers to examine a broader set of relationships and to test assumptions that were previously treated as analytically necessary simplifications. However, a caveat is necessary: increased volume does not automatically reduce omitted variable bias if the newly available variables are measured with error, are irrelevant, or if the data suffers from selection bias [2].

Velocity. This characteristic most significantly challenges the static interpretation of *ceteris paribus*. Classical approaches often imply that, at the moment of analysis, other relevant factors are held constant or treated as fixed. However, high-frequency and real-time data allow continuous monitoring of changing conditions. Variables are no longer assumed to remain unchanged but can be observed within their evolving dynamics. Thus, Velocity transforms *ceteris paribus* from a static mechanism of analytical isolation into a form of dynamic empirical control. In practical terms, this means that researchers can now assess whether the "other conditions" that are assumed to be constant in a model ***actually remain stable*** over the period of analysis - a test that was previously impossible with low-frequency data.

Variety. The integration of structured and unstructured data—including textual information, geolocation data, transactional records, and behavioral traces—expands the set of potentially relevant variables available for economic analysis. This broader informational environment can reduce the risk of omitted variable bias (OVB), one of the

principal limitations associated with the application of *ceteris paribus*. However, the inclusion of additional variables does not automatically resolve identification problems; theoretical guidance remains necessary for distinguishing relevant from irrelevant information. Moreover, unstructured data sources (e.g., text, images) introduce new measurement challenges that may offset the gains from increased variety [3].

Veracity. Data quality, completeness, consistency, and measurement reliability directly determine the extent to which "other conditions" can be empirically controlled. Low-quality or biased data may significantly reduce the analytical advantages created by increased volume and variety, demonstrating that the usefulness of Big Data depends not only on scale but also on reliability. This dimension is particularly important for *ceteris paribus* because measurement error in control variables can produce biased estimates of the relationship of interest - a problem that may worsen, rather than improve, with larger datasets [4].

Value. Ultimately, the methodological importance of Big Data is determined not by the quantity of information alone but by its capacity to improve causal interpretation, theoretical understanding, and decision-making quality. The value of Big Data lies in its ability to generate meaningful analytical insights and enhance the empirical assessment of economic relationships, rather than simply increasing the amount of available data. From the perspective of *ceteris paribus*, the "value" dimension asks a crucial question: Does the use of Big Data actually improve our ability to isolate causal effects, or does it merely increase the complexity of the analysis without improving identification?

The combination of these dimensions demonstrates that Big Data does not eliminate the logic of *ceteris paribus*; rather, it transforms how the principle is operationalized. The assumption of predetermined and artificially fixed "constant conditions" is gradually replaced by empirically grounded, data-driven mechanisms of control. This transition provides the foundation for the concept of the **empirical realization of *ceteris paribus* -** one that implies not the abstract isolation of variables, but the data-based monitoring and control of their dynamic interactions. However, this transition is neither automatic nor guaranteed. Big Data creates the ***potential*** for empirical operationalization of ***ceteris paribus***, but realizing this potential requires careful research design, appropriate methods for handling high-dimensional data, and theoretical guidance for distinguishing signal from noise.

In the era of Big Data, econometric analysis increasingly relies on large-scale, multidimensional, and high-frequency datasets. Such approaches enable researchers to incorporate a broader range of interconnected factors and to reduce their reliance on simplified assumptions and excessive abstractions that were often necessary under conditions of limited data availability.

Einav and Levin illustrate this shift empirically, showing how the growing availability of large-scale transactional and administrative data has expanded the scope of questions economists can address and reduced dependence on small, purpose-collected samples [5]. Unlike traditional empirical studies, which frequently relied on relatively small samples and predefined theoretical models, Big Data creates opportunities to draw on broader and more diverse sources of information. As a result, economic analysis becomes less dependent on simplified representations and increasingly grounded in the observation of actual behavioral patterns.

Varian [6] approaches this transformation from a methodological perspective, arguing that Big Data reshapes economic research primarily by giving researchers access to tools — drawn from computer science and statistics — capable of handling complexity and detail that traditional econometric methods were not designed to manage. However, the importance of Big Data is not determined solely by the volume of available information. Mayer-Schönberger and Cukier [7] extend this point from a more conceptual standpoint, emphasizing that Big Data changes the underlying logic of analysis itself: it facilitates the identification of complex relationships and patterns within large-scale datasets, including unstructured information, in ways that shift analysis from causal explanation toward correlational discovery.

From this perspective, Big Data reduces dependence on purely abstract modeling approaches and expands the possibilities for multidimensional analysis of economic processes. In conventional econometric models, researchers often specify a limited number of explanatory variables based on prior theoretical assumptions, whereas Big Data approaches enable the integration of a broader set of information and multiple interconnected features. As a result, models may achieve greater informational richness and reduce the risk of omitted variable bias.

Nevertheless, although Big Data substantially expands the analytical capabilities of economics, it does not automatically guarantee the correct identification of causal relationships. Large datasets may contain measurement errors, selection bias, irrelevant information, or relationships that are merely correlational rather than causal. Therefore, the effective use of Big Data requires the integration of economic theory, econometric methods, and advanced data analytics. As Athey and Imbens [8] emphasize, machine learning and Big Data methods are particularly powerful for prediction tasks; however, causal interpretation still requires theoretical foundations and appropriate identification strategies.

Big Data, then, does not replace *ceteris paribus* and does not transform economics into a purely data-driven discipline. Instead, it creates a new methodological environment in which theoretical abstraction and empirical evidence perform complementary functions. Within this hybrid framework, *ceteris paribus* evolves from a static assumption of isolation into a dynamically assessed and empirically informed analytical principle.

THE DIALECTICAL UNITY OF DEDUCTION AND INDUCTION

Traditional economic models have primarily been based on a deductive approach: researchers began with theoretical assumptions, developed formal models, and subsequently tested their consistency with empirical evidence. This approach can be summarized as the principle of **“assume and derive.”** Within this framework, theory determined which variables were considered relevant and how relationships among them should be interpreted.

The emergence of Big Data has significantly expanded this methodological framework by strengthening the role of empirical and inductive analysis. Contemporary technologies enable researchers to identify previously unobserved patterns, relationships, and behavioral regularities through the analysis of large-scale datasets. Consequently, the logic of inquiry is increasingly complemented by the principle of **“measure and explore,”**

whereby research begins not only with predefined theoretical hypotheses but also with systematic investigation of empirical regularities emerging from data.

This shift, however, should not be read as a linear transition from deduction to induction. Rather, it represents a dialectical interaction between these two modes of reasoning. In contemporary research, induction and deduction operate jointly and complement one another: empirical patterns identified through data analysis require theoretical interpretation and causal explanation, while theory provides the framework for determining which observed relationships carry economic significance and causal relevance. Big Data, then, does not diminish the role of theory; rather, it changes the nature of interaction between theoretical reasoning and empirical evidence.

A fuller understanding of this relationship requires the inclusion of **abductive reasoning** – a third form of inference that connects induction and deduction. Abduction begins when unexpected observations, anomalies, or previously unidentified patterns in the data prompt the formulation of new theoretical explanations, which are subsequently subjected to empirical testing and refinement. In Charles Sanders Peirce's formulation, abduction is the process of generating explanatory hypotheses to account for surprising or otherwise unexplained observations; it constitutes a mode of reasoning through which possible explanations are proposed and subsequently subjected to further inquiry [9]. Through abduction, the movement from **"data → theory → data"** is transformed from a linear sequence into an open-ended iterative process in which each cycle contributes to the refinement of both theoretical frameworks and empirical understanding.

The connection between abduction and *ceteris paribus* is particularly significant. When unexpected patterns or anomalies emerge in Big Data analysis, they often indicate that the *ceteris paribus* condition has been violated—that is, some factor that was assumed to be "held constant" has actually varied, or some unmodeled interaction has occurred. Abductive reasoning directs researchers to identify these violations and incorporate them into revised theoretical frameworks. In this sense, the dialectical spiral is not merely a methodological preference but a response to the empirical failures of static *ceteris paribus* assumptions.

Figure 1 illustrates this dialectical process. The spiral begins with theory (deduction) and moves through empirical observation (induction) to the identification of anomalies or unexpected patterns (abduction), which then generate new theoretical questions and empirical tests. At each cycle, the *ceteris paribus* assumption is reassessed and refined: variables that were previously abstracted may be brought back into the model, and variables that were previously included may be shown to be irrelevant under certain conditions. **The CPV coefficient, introduced in Section 12, provides a formal measure of how much empirical evidence supports the *ceteris paribus* assumptions at each stage of this spiral.**

Therefore, in the Big Data era, economic knowledge formation can be understood as a dynamic dialectical process involving the continuous interaction of deduction, induction, and abduction. This perspective provides the methodological foundation for a hybrid paradigm in which theoretical abstraction and empirical discovery are not competing approaches, but mutually reinforcing components of economic analysis.

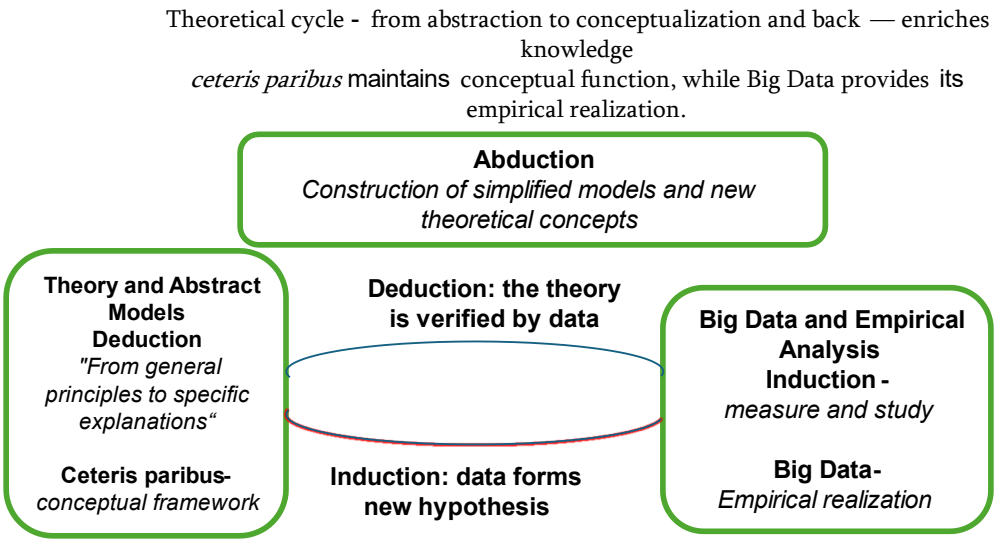


Figure 1: The Dialectical Spiral of Deduction, Induction, and Abduction: From *Ceteris Paribus* Abstraction to Big Data-Based Empirical Realization

The contemporary literature on Big Data and machine learning likewise highlights this transformation: Big Data expands the capabilities of econometric analysis by providing access to larger, more granular, and more diverse datasets, thereby enabling economic processes to be examined in greater detail [6, 10].

MACHINE LEARNING AS A TOOL FOR EXTENDING THE *CETERIS PARIBUS* PRINCIPLE

Machine Learning (ML) methods not only complement traditional econometrics but also fundamentally transform the functions that the *ceteris paribus* assumption has traditionally performed.

Machine Learning as a Variable Selection Tool

The classical econometric approach requires the researcher to determine in advance which variables are "important" and which should be abstracted away under the *ceteris paribus* assumption. This decision is inherently subjective. Machine Learning (ML) methods make it possible to analyze a large number of potential variables simultaneously and assess their relevance, thereby reducing the risk that important variables will be excluded from the model due to the researcher's subjective judgment. In particular, regularization-based methods, such as Lasso (L1), Ridge (L2), and Elastic Net, as well as tree-based algorithms, facilitate feature selection while controlling model complexity and mitigating overfitting.

From the perspective of the CPV, ML-based variable selection provides a quantitative assessment of the $N_{modeled}/N_{relevant}$ ratio, while variables identified as “important” by the algorithm become empirically informed candidates for inclusion in the model.

Table 1: Comparison of ML Variable Selection Methods.

Method	Principle	Variable Selection	Relationship with *ceteris paribus*
Lasso	L1 regularization	Performs variable selection (some coefficients are reduced to zero)	Reduces $N_{relevant}$ — "determines" which variables are truly important
Ridge	L2 regularization	Does not perform variable selection (all variables remain in the model)	Preserves the idea of *ceteris paribus*, but reduces the influence of less important variables
Elastic Net	Combination of Lasso + Ridge	Performs variable selection	Combines the advantages of both Lasso and Ridge methods
Random Forest	Ensemble of decision trees	Performs variable selection (feature importance analysis)	Identifies the variables with the greatest influence

However, a crucial caveat is that variable importance in ML models does not imply causal relevance. Variables that are strong predictors of Y may not be causal determinants, and variables that are weak predictors may still have causal effects that are masked by confounding or measurement error [11]. Therefore, ML-based variable selection is a *starting point* for identifying candidate variables, not a substitute for theoretical reasoning about causal structure.

Machine Learning as a Tool for Detecting Interactions

The *ceteris paribus* assumption essentially implies that variables operate independently of one another. In real-world economies, however, this condition is rarely satisfied. ML methods, such as Random Forest (which incorporates feature interactions when partitioning the feature space), XGBoost (gradient boosting), and neural networks, automatically detect complex, nonlinear interactions among variables. This changes the logic of *ceteris paribus*: instead of assuming that "all other conditions remain constant," ML reveals how the effect of one variable changes depending on the values of other variables.

Table 2: Comparison of Interaction Modeling in Classical Econometric and Machine Learning Approaches.

Dimension	Classical Model	ML Model
Interactions	Pre-specified (e.g., $X_1 \times X_2$)	Complex interactions automatically detected from data
Example	"Education × Gender"	"Education × Gender × Age × Region"
Ceteris paribus	"Other conditions remain constant"	Interactions are explicitly revealed and modeled
Outcome	Simple and interpretable, but potentially less accurate	More complex model, often with higher predictive performance

Machine Learning as a Prediction Tool

Models based on *ceteris paribus* often provide strong explanations but weak predictions. Machine Learning shifts the emphasis from explanation toward prediction. As Mullainathan

and Spiess note, the question “Which variables are important?” is replaced by the question “How can we achieve the most accurate prediction?” [12].

Tree-based methods, particularly Random Forest and XGBoost, partition the data space into homogeneous regions; neural networks learn complex nonlinear relationships; and Support Vector Machines (SVMs) identify optimal decision boundaries. Together, these methods are particularly well suited for high-dimensional prediction tasks. In relation to the *ceteris paribus* assumption, this shift can be expressed as follows: whereas *ceteris paribus* asks, “What happens if only X changes while all other factors remain constant?”, machine learning asks, “What happens when many variables change simultaneously and interact with one another?”

This shift from explanation to prediction has important implications for the role of *ceteris paribus*. In a purely predictive context, the question "What happens if X changes while all other factors remain constant?" is less relevant than the question "What is the most accurate prediction given the current values of all variables?" However, for policy purposes, the *ceteris paribus* question remains central: policymakers need to know the causal effect of changing X, not just the predictive relationship. This is why hybrid approaches that combine ML prediction with structural identification are increasingly important [8].

Machine Learning as a Support Tool for Causal Inference

Although ML methods such as Random Forest and neural networks do not, by themselves, provide causal interpretations, they strengthen the tools of causal analysis. Athey and Imbens discuss how ML is applied within a counterfactual analysis framework, where algorithms help identify complex relationships and structures present in high-dimensional data [8].

Causal Forest determines for which individuals a particular policy intervention is effective [13], while Double/Debiased Machine Learning applies ML to instrumental variable estimation and high-dimensional control settings [14]. These methods do not replace the *ceteris paribus* principle; rather, they provide quantitative tools for assessing the conditions under which its assumptions hold. In particular, they reveal how causal effects vary across different populations through the estimation of heterogeneous treatment effects, a dimension that the traditional *ceteris paribus* framework, which typically focuses on a single average treatment effect, is unable to capture.

The “Black Box” Problem and Interpretable Machine Learning

One of the main criticisms of ML models is related to their lack of interpretability. They are often perceived as “black boxes.” The *ceteris paribus* approach, in contrast, inherently requires interpretability, as it addresses the question: “What happens if X changes?” To resolve this tension, Interpretable ML methods have been developed, which restore the logic of *ceteris paribus* within ML models.

PDP and ICE methods provide an ML-based interpretation of the *ceteris paribus* principle: they show how model predictions change when a specific variable is varied while other variables are held constant, thereby offering an empirically testable approximation of the assumption of “other things being equal.” Thus, while they do not eliminate the

"black box" problem entirely, they provide an important methodological tool for improving the interpretability of complex machine learning models.

Table 3: Comparison of Interpretable Machine Learning Methods (LIME, SHAP, PDP, ICE) and Their Relationship with Ceteris Paribus.

Method	How It Works	Relationship with *ceteris paribus*
LIME	Builds a local model that explains a specific prediction	Answers the question: "What happens if X changes?" (similar to the *ceteris paribus* logic)
SHAP	Uses Shapley Values to determine the contribution of each variable to a specific prediction	Answers the question: "What is the contribution of each variable?" (a generalization of *ceteris paribus*)
PDP (Partial Dependence Plot)	Shows, on average, how Y changes when X changes, taking into account the distribution of other variables	Represents an approximate empirical form of the *ceteris paribus* idea
ICE (Individual Conditional Expectation)	Similar to PDP, but analyzes changes at the individual level	Represents an individual-level version of *ceteris paribus*

The importance of interpretable ML for *ceteris paribus* is that these methods provide an empirical approximation of what the theoretical assumption would imply. PDP and ICE show how predictions change when a variable is varied; if the relationship is nonlinear or interacts with other variables, this becomes immediately visible. This provides researchers with diagnostic information: if the PDP shows that the effect of X on Y changes dramatically depending on the values of Z, then the *ceteris paribus* assumption that "all other things are equal" may be untenable. The researcher can then revise the model to incorporate the interaction, moving from a static abstraction to a more empirically grounded specification.

ARTIFICIAL INTELLIGENCE (AI) AND LARGE LANGUAGE MODELS: BEYOND MACHINE LEARNING

Artificial Intelligence (AI) is not merely an extension of Machine Learning; it fundamentally changes the logic of analysis. In particular, the development of Large Language Models (LLMs) has expanded the possibilities for data analysis, knowledge organization, and hypothesis generation.

Three characteristics distinguish AI from traditional ML:

- **First, generativity.** AI systems are capable not only of classifying, analyzing, or predicting based on existing data but also of generating new data, hypotheses, and models.
- **Second, natural language interface.** Researchers can "communicate" with data through natural language, enabling more intuitive interaction with analytical systems.
- **Third, automated model construction.** AI can automatically construct models, test hypotheses, and derive conclusions.

It is precisely through these dimensions that AI transforms the logic of research. While ML addresses the question: **"What is the effect of X on Y based on available data?"**, AI can go further by asking: **"What could be the effect of X on Y?"**, **"What possible mechanisms could explain how X affects Y?"**, and **"How can these mechanisms be empirically tested?"**.

AI as an Automatic Hypothesis Generator

As discussed in Section 4, abductive reasoning (inference to the best explanation) provides a methodological link between induction and deduction. Artificial intelligence (AI), particularly Large Language Models (LLMs), extends this form of reasoning by enabling the systematic exploration of a broader space of possible explanations and hypotheses. Rather than requiring researchers to determine entirely *ex ante* which variables are theoretically important ($N_{relevant}$), AI-based approaches can assist in generating candidate variables, discovering hidden relationships, and expanding the set of factors considered in model specification. This has direct implications for the components of the CPV framework: by increasing the number of theoretically and empirically relevant variables incorporated into the model, AI may increase the $N_{modeled}/N_{relevant}$ ratio. At the same time, it may reduce I_{abs} (the index of influence of abstracted variables) by identifying factors that, under a traditional *ceteris paribus* approach, would often remain excluded from the analytical framework.

Explainable AI (XAI): From Abstraction to Empirical Results

The logic of interpretable ML discussed in Section 5 is extended by Artificial Intelligence to complex, multi-layered systems through Explainable AI (XAI). XAI methods (LIME, SHAP, and Integrated Gradients) demonstrate how each variable contributes to the outcomes generated by AI models.

The particular importance of XAI in the context of *ceteris paribus* lies in the fact that it changes the epistemological status of the assumption: *ceteris paribus* is no longer an **a priori** assumption, but becomes an empirical outcome. AI demonstrates which variables can be considered "constant" and how significant each variable is in determining the result [15, 16].

AI as a Tool for Hybrid Modeling

The most significant contribution of AI to economic modeling lies in its ability to enable **hybrid modeling**. AI can integrate four analytical layers within a single coherent framework:

- **Economic theory**: fundamental relationships (e.g., AD–AS, IS–LM) that provide causal structure and ensure interpretability;
- **Big Data:** real-time, high-frequency empirical information;
- **Machine Learning:** nonlinear models capable of capturing complex interactions among variables;

- **AI-driven automated analysis:** hypothesis generation, automated model construction, and inference.

The relationship with *ceteris paribus* becomes particularly evident in this context: hybrid modeling combines the abstraction provided by *ceteris paribus* with the empirical richness of Big Data, while AI determines which variables should be abstracted (within the *ceteris paribus* framework) and which should be incorporated into the analysis through data-driven approaches.

This constitutes the operational logic underlying the CPV coefficient, which aims to formalize the relationship between theoretical abstraction and empirical representation, thereby providing the conceptual foundation for the proposed hybrid paradigm.

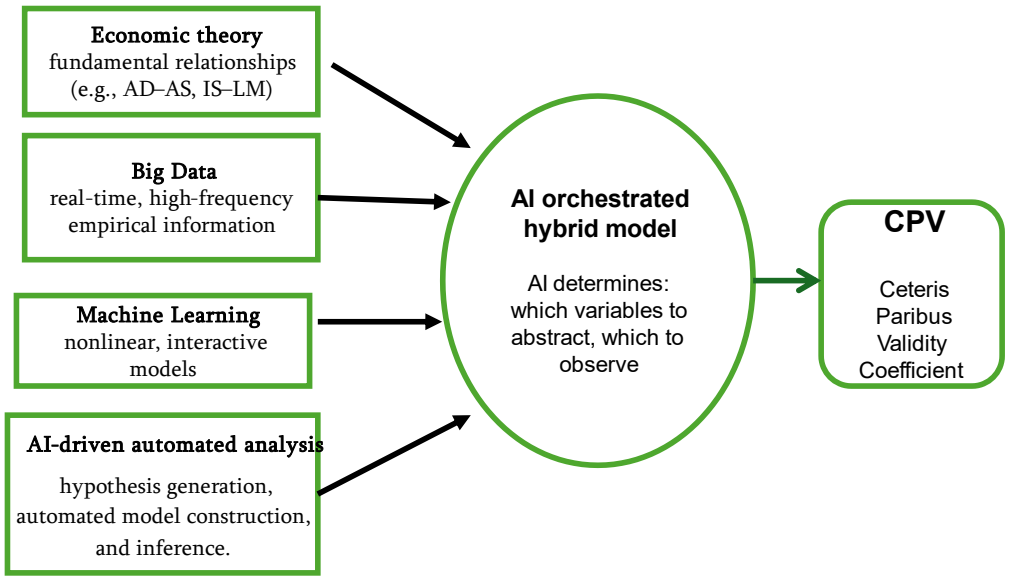


Figure 2: AI-Orchestrated Hybrid Model: Integration of Theory, Big Data, Machine Learning, and AI-Driven Automated Analysis for Assessing the Validity of Ceteris Paribus (CPV).

At the same time, the use of AI and LLMs introduces its own methodological and ethical risks, including instability and limited reproducibility of results, the generation of hypotheses that may lack empirical validation, and dependence on training data. These challenges further reinforce the need for a robust theoretical and normative framework to guide the responsible integration of AI into economic modeling.

Limitations of AI in Economic Modeling

Despite its potential, the integration of AI into economic modeling introduces significant challenges that require explicit acknowledgment.

First, instability and reproducibility. AI-generated results can be sensitive to training data, model initialization, and hyperparameter choices. Different runs may produce different hypotheses or conclusions, raising concerns about the reliability of AI-assisted analysis [17].

Second, hallucination and generated false hypotheses. LLMs may generate hypotheses that appear plausible but are empirically unfounded or logically inconsistent. Unlike traditional hypothesis generation, which is constrained by theoretical reasoning, AI models may produce "creative" but untestable claims.

Third, dependence on training data. AI models learn from historical data, including its biases and limitations. If the training data reflects historical inequalities or measurement errors, AI-generated hypotheses may reproduce or amplify these problems.

Fourth, lack of causal understanding. AI models are pattern-matching systems; they do not possess an understanding of causal mechanisms. This limitation is particularly important in economics, where understanding causal structure is central to policy analysis.

These challenges reinforce the argument that AI should be used as a *complement* to, not a replacement for, economic theory and rigorous research design. The hybrid paradigm proposed in this paper explicitly acknowledges this limitation by insisting on theoretical grounding and causal structure as the foundation of economic modeling.

AGENT HETEROGENEITY AND AGENT-BASED MODELING

Traditional economic models have largely relied on the concept of the **representative agent**, according to which the behavior of a single, typical individual is assumed to represent the behavior of all economic agents. While this abstraction greatly simplifies analytical modeling, it inevitably suppresses heterogeneity across individuals in terms of income, preferences, expectations, behavioral characteristics, and decision-making strategies.

The emergence of **Big Data** has fundamentally challenged this paradigm. The availability of large-scale, high-frequency, and highly disaggregated datasets allows researchers to observe the behavior of millions of individuals, firms, and institutions rather than relying on aggregate averages. Consequently, modern empirical research increasingly focuses on **heterogeneous effects**, identifying how policies and economic shocks affect different population groups, regions, sectors, or income classes. For example, monetary policy may produce substantially different responses across households with varying wealth, liquidity constraints, or expectations.

The growing recognition of heterogeneity has also stimulated the development of **Agent-Based Modeling (ABM)**. Unlike representative-agent models or standard equilibrium frameworks, ABM explicitly models each economic actor (households, firms, banks, or governments) as an autonomous agent with its own objectives, information, adaptive behavior, and decision rules. Aggregate macroeconomic outcomes then emerge endogenously from the interactions among heterogeneous agents through a **bottom-up** process rather than being derived from predetermined aggregate assumptions.

Nevertheless, ABM should not be viewed as a universal substitute for conventional economic modeling. Several methodological limitations remain. First, ABMs are computationally intensive and require substantial computational resources together with detailed micro-level data. Second, simulation results are often highly sensitive to the specification of behavioral rules, parameter values, and initial conditions; relatively small modifications may generate substantially different macroeconomic outcomes. Third, although ABMs successfully demonstrate how aggregate phenomena emerge from

decentralized interactions, they do not, by themselves, provide rigorous causal identification. Their primary explanatory contribution lies in generating plausible mechanisms and exploring complex dynamics rather than establishing causal effects through statistical inference. Recent research has therefore increasingly focused on combining empirical data with machine learning techniques to estimate and calibrate behavioral rules more systematically, thereby reducing reliance on ad hoc assumptions [18].

These limitations suggest that the greatest methodological potential lies in **hybrid modeling**, where ABM is integrated with econometric methods, machine learning, and Big Data analytics. Such integration allows researchers to adjust behavioral parameters using real-world microdata, validate simulation outcomes through econometric analysis, and improve predictive performance while maintaining theoretical interpretability. This emerging hybrid paradigm combines the strengths of simulation-based modeling, causal inference, and data-driven learning.

A well-known example is the analysis of financial crises. Farmer and Foley [19] demonstrate that systemic risk may emerge endogenously from interactions among heterogeneous agents, challenging the representative-agent paradigm. However, the credibility of such simulations depends on empirical and econometric validation. Consequently, combining ABM with econometric techniques produces more reliable and policy-relevant results than either methodology alone. ABM contributes behavioral realism and generative explanations, whereas econometrics provides causal discipline, statistical inference, and empirical validation. Overall, Big Data has shifted economic analysis from representative-agent abstractions toward heterogeneous, micro-founded approaches. Yet this transformation does not eliminate the role of abstract theoretical models. Instead, it supports a new methodological paradigm in which theory, Big Data, machine learning, econometrics, and agent-based simulation operate as complementary components of a unified hybrid framework for understanding increasingly complex economic systems [2, 18].

The relationship between ABM and *ceteris paribus* is subtle and deserves explicit attention. In a traditional representative-agent model, the *ceteris paribus* assumption is applied at the aggregate level: we ask how a policy change affects the representative agent, holding other factors constant. In an ABM framework, this logic is extended to the micro level: we ask how a policy change affects *different types of agents* under *different contextual conditions*. The *ceteris paribus* assumption becomes conditional: "holding other factors constant, what is the effect for agent type A?" This conditional approach is more realistic but also more complex to implement and validate. The CPV coefficient can be extended to capture this: CPV(ABM) would measure the extent to which model outcomes are robust to alternative agent specifications, behavioral rules, and initial conditions.

REAL-TIME ECONOMICS AND NOWCASTING

Beyond enabling a more detailed representation of agent heterogeneity, **Big Data** has introduced another fundamental transformation in economic analysis: the ability to assess economic activity in **near real time**, giving rise to the concept of the **Real-Time Economy**. Traditional economic models were largely based on a static analytical framework, relying on periodically collected and often delayed official statistics. As a result, such models primarily provided economic conditions after they had occurred, limiting their ability to

respond rapidly to emerging changes. In contrast, the Big Data environment provides continuously updated information streams that allow researchers and policymakers to monitor economic dynamics in real time and adjust analytical models as new information becomes available.

A key manifestation of this transformation is the development of **nowcasting**—the estimation of current economic conditions before official statistics are released by exploiting high-frequency data sources, including digital transactions, online activity indicators, mobility data, and other alternative datasets. The GDP nowcasting framework developed by **Giannone, Reichlin, and Small** integrates a large number of economic indicators observed at different frequencies within a dynamic factor model, enabling timely estimates of current economic activity despite incomplete or delayed official data releases [20].

More recent advances have further expanded nowcasting through the integration of machine learning methods with high-dimensional and large-scale datasets, allowing researchers to exploit nonlinear relationships, improve predictive performance, and develop more flexible forecasting frameworks [18].

From the perspective of **ceteris paribus**, the methodological significance of nowcasting lies in its challenge to the traditional static interpretation of the assumption. Rather than treating other relevant factors as fixed, nowcasting continuously observes, updates, and incorporates changes in these factors through real-time information flows. In this sense, it transforms the traditional assumption of “holding other things constant” into a framework of **continuous empirical monitoring and dynamic adjustment**. Consequently, the role of ceteris paribus as a static abstraction is reduced, while dynamic empirical control becomes increasingly important in economic modeling.

This development is directly connected to the extended Ceteris Paribus Validity (CPV) framework proposed in this study. Specifically, nowcasting enhances the temporal dimension of empirical validity by increasing the model's capacity to incorporate newly available information over time. In the CPV framework, nowcasting improves the "time-related component" of the Big Data Coverage Coefficient: it reduces the delay between data collection and model updating, allowing the model to better represent the "true" state of the economy at the time of analysis. By improving temporal coverage and reducing information delays, nowcasting strengthens the empirical validity of economic models and increases their capacity to represent rapidly changing economic environments. However, the Lucas critique applies to nowcasting as well: if the structural relationships governing the economy change, even real-time data may not provide reliable guidance.

CORRELATION, CAUSATION, AND CAUSAL INFERENCE

Although Big Data has expanded the empirical scope of economic analysis, increased data availability does not eliminate the methodological challenge of identifying causal effects. Causal interpretation continues to depend on theoretical assumptions, identification strategies, and empirical validation. By allowing researchers to analyze many variables simultaneously, Big Data can reveal complex interactions and previously hidden patterns. Nevertheless, without an appropriate causal framework, these patterns may remain

descriptive rather than explanatory. This limitation is not merely theoretical; it is illustrated by several empirical examples:

- **Google Flu Trends:** A Big Data-based forecasting system initially identified a strong correlation between Google search queries and influenza prevalence. However, its predictions later failed systematically because the model became excessively optimized to historical data patterns and did not adequately capture the underlying causal mechanisms of influenza transmission. This case demonstrates the risks of relying exclusively on predictive correlations without sufficient theoretical foundations and causal validation [21].
- **Ice cream sales and drowning incidents:** A strong statistical correlation exists between these two variables, but there is no direct causal relationship between them. Instead, both are driven by a common underlying factor, namely temperature and seasonal conditions. This classic example illustrates the importance of identifying **confounding variables**, which may generate misleading associations between observed variables.
- **Social media data:** Large-scale behavioral data from social networks provide valuable descriptive insights into the diffusion of opinions, preferences, and trends. However, causal interpretation remains limited due to challenges such as selection bias, self-selection mechanisms, and the non-random nature of online participation. Consequently, observed behavioral patterns may not necessarily represent causal effects in the broader population.

Therefore, while Big Data provides unprecedented opportunities for discovering patterns and improving predictive performance, establishing causal relationships requires additional methodological frameworks, including experimental designs, econometric identification strategies, and theory-based modeling. The central challenge of modern economic analysis is not merely to identify statistical regularities within large datasets, but to understand the mechanisms through which economic relationships emerge.

These examples demonstrate why **causal inference—and, more broadly, the logic of ceteris paribus—remains essential even in the Big Data era**. Correlation describes statistical associations between variables, whereas causal analysis seeks to determine what would change if a specific factor were altered while other relevant conditions were controlled. In this sense, causal inference represents **an empirical reconstruction** of the ceteris paribus principle: rather than assuming other factors remain constant purely through theoretical abstraction, modern econometric methods attempt to approximate this condition through identification strategies, experimental designs, and statistical controls.

Thus, Big Data does not eliminate the need for ceteris paribus reasoning; instead, it transforms its role. The traditional abstract assumption of holding other things constant is increasingly complemented, and partially operationalized, through empirical methods capable of isolating causal effects within complex economic systems. This transformation is bidirectional: causal inference methods provide the empirical discipline that makes the ***ceteris paribus*** assumption testable, while Big Data provides the empirical foundation for implementing these methods with greater precision and scope.

The importance of this methodological direction was particularly highlighted by the 2021 Nobel Prize in Economic Sciences, awarded to David Card, Joshua Angrist, and Guido

Imbens for their contributions to empirical research and causal inference. Their methodological contributions have direct relevance for the interpretation and application of the *ceteris paribus* principle and can be understood through two main dimensions.

First, Angrist and Imbens developed the concept of the Local Average Treatment Effect (LATE), which extends the traditional interpretation of *ceteris paribus* beyond a single average treatment effect by introducing the question: **"For whom does the effect apply?"** and recognizing that treatment effects may differ across individuals. Rather than assuming a universal effect for all agents, LATE identifies the causal effect for a specific subgroup, known as compliers, defined by their response to the treatment assignment mechanism [22]. This is a crucial refinement: in a heterogeneous world, the ***ceteris paribus*** assumption must be conditional on the population under study.

Second, Card's contribution through the use of natural experiments demonstrates how naturally occurring quasi-random variation in real-world settings can be used to approximate the condition of "holding other things constant." [23, 24]. By comparing otherwise similar groups exposed to different conditions, such as groups affected differently by a policy change, natural experiments create a more credible counterfactual framework for estimating causal effects. In the CPV framework, the validity of such natural experiments depends on how well the comparison groups approximate the ***ceteris paribus*** condition, a question that can be empirically assessed using the data richness provided by Big Data.

Together, these methodological advances transform the *ceteris paribus* principle from a purely theoretical assumption into an empirically testable framework for identifying causal effects under realistic conditions. Causal inference methods do not replace the *ceteris paribus* principle; rather, they strengthen its empirical foundation by providing tools for evaluating when and under what conditions causal assumptions hold.

Although Big Data and the growing availability of large-scale datasets substantially expand the possibilities for economic analysis, theoretical reasoning and research design remain essential for determining whether observed relationships represent genuine causal effects. In this sense, modern causal inference provides a methodological bridge between theoretical abstraction and empirical validation by translating causal assumptions into explicit identification strategies and counterfactual comparisons [25, 26].

DEPENDENCE ON HISTORICAL DATA: THE LUCAS CRITIQUE AND OVERFITTING

One of the important limitations of **Big Data** is its dependence on historical data. Algorithms primarily learn patterns and regularities observed in the past; however, economic agents continuously adapt to changes in policy regimes, technological innovations, crises, and shifts in expectations. Two seemingly distinct manifestations of this problem, the **Lucas critique** and **overfitting,** actually originate from a common underlying challenge and should therefore be considered together.

The Lucas critique highlights a fundamental challenge for empirical and data-driven modeling: historical regularities cannot always be assumed to remain valid when economic structures change. As Lucas demonstrated, changes in policy regimes and agents' expectations can alter the behavioral relationships underlying observed data [27]. Therefore, reliable economic modeling requires not only predictive accuracy but also theoretical understanding of structural mechanisms and behavioral responses.

The problem of **overfitting** arises when a model becomes excessively tailored to specific patterns, including random fluctuations, contained in historical data. Such models may reproduce past observations with high accuracy, but lose predictive power when applied to new environments or future conditions. In machine learning, this represents a fundamental trade-off between fitting existing data and achieving generalization beyond the training sample.

Both problems reflect a common methodological challenge: historical patterns, whether behavioral relationships or statistical regularities, may not remain stable over time. Therefore, the central question is how to construct models that exploit the informational advantages of Big Data while avoiding excessive dependence on historical patterns.

One promising solution is the integration of **structural economic models** (such as Dynamic Stochastic General Equilibrium—DSGE models) with Big Data approaches. Structural frameworks provide theoretically grounded relationships and relatively stable behavioral parameters, while Big Data and machine learning methods enable continuous empirical calibration, parameter updating, and improved forecasting performance. From this perspective, economic abstraction, despite its simplifying nature, may sometimes yield more stable and generalizable results because it focuses on fundamental relationships.

ETHICAL AND INSTITUTIONAL CHALLENGES

The use of **Big Data** is accompanied by significant ethical and institutional challenges. The increasing reliance on algorithm-based decision-making raises fundamental questions regarding accountability, transparency, and potential biases embedded in data and analytical systems. If historical datasets reflect existing social or economic inequalities, algorithms may not only reproduce these biases, but also reinforce them through automated decision processes. This issue further demonstrates that the availability of massive amounts of data cannot substitute for appropriate theoretical, ethical, and normative frameworks.

Therefore, the central issue is not a confrontation between **Big Data** and traditional economic methodologies, as both approaches possess distinct advantages and limitations. Big Data expands the empirical foundation of economic analysis, reveals complex relationships, and enhances predictive capabilities; however, it does not, by itself, provide a complete explanation of causal mechanisms. Economic theory, abstraction, and the **ceteris paribus** principle, in contrast, provide essential tools for interpreting economic processes and understanding causal relationships, although their simplifying assumptions may fail to capture the full complexity of real-world systems.

This distinction highlights the fundamental difference between **prediction and explanation**: Big Data approaches often provide answers to the question **"What will happen?"**, whereas economic theory seeks to answer the deeper question **"Why does it happen?"**. Effective economic analysis therefore requires not only predictive accuracy, but also theoretical interpretation and causal understanding.

Future economic models will likely be based on a **hybrid paradigm** that integrates theoretical reasoning, large-scale data analysis, machine learning techniques, and causal inference methods [12]. Such an approach is particularly important in a modern economic environment where agents can no longer be understood solely as fully rational and

homogeneous entities. Instead, economic behavior is shaped by information limitations, institutional constraints, behavioral factors, and heterogeneous preferences [28, 29].

TOWARD A CETERIS PARIBUS VALIDITY (CPV) FRAMEWORK

The dialectical framework developed above implies that *ceteris paribus* assumptions should be continuously reassessed against empirical evidence. This raises the question: Can such reassessment be quantified? This section briefly outlines the **Ceteris Paribus Validity (CPV) framework** as a conceptual instrument for this purpose. Its full formalization and empirical application are the subject of ongoing research.

The CPV framework is motivated by two intuitive dimensions: **coverage** (the proportion of potentially relevant variables included in the model) and **control** (the estimated influence of excluded variables on the relationship of interest). Higher CPV values indicate stronger empirical grounding of the *ceteris paribus* assumption.

Two extensions are particularly relevant in the Big Data context. **CPV(BD)** incorporates data quality (Veracity), temporal resolution (Velocity), and informational breadth (Variety). **AI-CPV** extends the framework to incorporate AI-driven hypothesis generation and automated model validation.

The CPV framework does not replace existing evaluation criteria (R^2, AIC, cross-validation) but complements them by addressing a distinct question: not only "Does the model predict well?" but also "Are its *ceteris paribus* assumptions empirically defensible?"

Implementation faces challenges, including the difficulty of defining "relevant" variables and estimating the influence of excluded factors. The CPV is therefore best understood as a diagnostic heuristic that can guide research design and interpretation, rather than as a definitive test.

Formalization, estimation strategies, and empirical applications are reserved for future work. In the present article, the CPV framework serves to operationalize the dialectical logic developed in the preceding sections.

CONCLUSION

The transformation of economic modeling in the data era can be identified along three major dimensions that have been developed throughout this paper.

First, **there is a shift from the artificial assumption of "other factors remaining constant" toward their empirical, data-driven control.** Big Data analytical methods, including **nowcasting** and **Agent-Based Modeling**, make it possible to move beyond treating variables as "frozen" in a static framework and instead monitor their evolution dynamically. In this sense, the traditional role of ceteris paribus as a form of **static isolation** is increasingly replaced by **dynamic empirical control**.

Second**, the number of variables incorporated into economic analysis has expanded substantially**. Instead of focusing on a limited set of key factors, modern data-driven approaches allow researchers to consider thousands of interconnected variables, thereby reducing the risk of **omitted variable bias (OVB)** and providing a more comprehensive representation of complex economic relationships.

Third, **the logic of scientific inquiry itself is changing.** While traditional economic analysis was primarily deductive—moving from theoretical assumptions and hypotheses toward empirical testing—the Big Data environment strengthens the role of **induction** through large-scale data analysis and **abduction** through the generation of new hypotheses from unexpected patterns, anomalies, and empirical regularities. These processes become important sources for extending and refining economic theory.

These three transformations create a **dialectical spiral cycle** in which abstraction (deduction) and empirical reality (induction) interact through abduction, with each cycle enriching and improving the next stage of knowledge development. From this perspective, ceteris paribus is no longer an a priori assumption fixed once and for all; rather, it becomes a dynamic component of the research process and a conceptual instrument that is continuously reassessed in light of empirical evidence. The CPV coefficients provide a formal metric for this reassessment: they quantify the extent to which empirical evidence supports the "other things equal" assumption at each stage of the analytical cycle.

This interpretation also implies a reconsideration of **Friedman's instrumentalist view of economic methodology**: the value of a model should not be evaluated solely by its predictive accuracy, but also by the plausibility and relevance of its assumptions within a complex, multi-factor, data-driven environment. A model that relies on strong *ceteris paribus* assumptions that are empirically unsupported may yield accurate short-run predictions, but remain fragile under structural change, a lesson illustrated by failures such as Google Flu Trends.

Ultimately, the paradigm of **hybrid modeling**, which integrates economic theory (providing causal structure), Big Data (providing a broad empirical foundation), and AI/Machine Learning (capturing nonlinear relationships, interactions, and predictive patterns), offers an integrated framework that not only preserves the methodological value of ceteris paribus, but also strengthens it in the Big Data era. This paradigm creates new opportunities for a more comprehensive, dynamic, and causally grounded analysis of economic processes. Such an approach is particularly relevant in the contemporary digital economy, where agents can no longer be understood as fully rational and homogeneous entities, and where data has become a central resource for economic decision-making.

Future research should focus on three directions: (1) empirical applications of the CPV coefficients to specific economic models, assessing their diagnostic value; (2) development of computational tools for automating CPV estimation using ML and AI methods; and (3) integration of the CPV framework with existing model evaluation criteria (e.g., AIC, BIC, cross-validation) to provide a more complete assessment of model validity. These extensions will further strengthen the methodological foundations of the hybrid paradigm proposed in this paper.

DECLARATIONS

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflicts of interest.

Data Availability: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Author Contributions: Conceptualization, M.T. and A.G.; writing – original draft, M.T. and A.G.; writing – review and editing, M.T. and A.G.

Ethics/Consent: Not applicable. This study did not involve human participants or animals.

Use of Generative AI: The article was originally written in Georgian. A generative AI tool (Claude, Anthropic) was used to translate and language-edit the manuscript draft. The authors checked the content, sources, and conclusions and accept full responsibility for them.

REFERENCES

[1]. Zikopoulos, P., C. Eaton, D. deRoos, T. Deutsch, and G. Lapis, *Understanding big data: Analytics for enterprise class Hadoop and streaming data*. 2012, New York: McGraw-Hill.

[2]. Athey, S., *Beyond prediction: Using big data for policy problems.* Science, 2017. 355(6324): p. 483-485. DOI: https://doi.org/10.1126/science.aal4321.

[3]. Gentzkow, M., B.T. Kelly, and M. Taddy, *Text as data.* Journal of Economic Literature, 2019. 57(3): p. 535-574. DOI: https://doi.org/10.1257/jel.20181020.

[4]. Bound, J., C. Brown, and N. Mathiowetz, *Measurement error in survey data*, in *Handbook of econometrics*, J.J. Heckman and E. Leamer, Editors. 2001, Elsevier: Amsterdam. Vol. 5, p. 3705-3843. DOI: https://doi.org/10.1016/S1573-4412(01)05012-7.

[5]. Einav, L. and J. Levin, *The data revolution and economic analysis*, in *Innovation policy and the economy*, J. Lerner and S. Stern, Editors. 2014, University of Chicago Press: Chicago. Vol. 14, p. 1-24. DOI: https://doi.org/10.1086/674019.

[6]. Varian, H.R., *Big data: New tricks for econometrics.* Journal of Economic Perspectives, 2014. 28(2): p. 3-28. DOI: https://doi.org/10.1257/jep.28.2.3.

[7]. Mayer-Schönberger, V. and K. Cukier, *Big data: A revolution that will transform how we live, work, and think*. 2013, Boston: Houghton Mifflin Harcourt.

[8]. Athey, S. and G.W. Imbens, *Machine learning methods that economists should know about*. Annual Review of Economics, 2019. 11: p. 685-725. DOI: https://doi.org/10.1146/annurev-economics-080217-053433.

[9]. Peirce, C.S., *Collected papers of Charles Sanders Peirce*. Vol. 5. 1934, Cambridge, MA: Harvard University Press.

[10]. Einav, L. and J. Levin, *Economics in the age of big data.* Science, 2014. 346(6210): p. 1243089. DOI: https://doi.org/10.1126/science.1243089.

[11]. Hastie, T., R. Tibshirani, and J. Friedman, *The elements of statistical learning: Data mining, inference, and prediction*. 2nd ed. 2009, New York: Springer. DOI: https://doi.org/10.1007/978-0-387-84858-7.

[12]. Mullainathan, S. and J. Spiess, *Machine learning: An applied econometric approach.* Journal of Economic Perspectives, 2017. 31(2): p. 87-106. DOI: https://doi.org/10.1257/jep.31.2.87.

[13]. Wager, S. and S. Athey, *Estimation and inference of heterogeneous treatment effects using random forests.* Journal of the American Statistical Association, 2018. 113(523): p. 1228-1242. DOI: https://doi.org/10.1080/01621459.2017.1319839.

[14]. Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins, *Double/debiased machine learning for treatment and structural parameters.* The Econometrics Journal, 2018. 21(1): p. C1-C68. DOI: https://doi.org/10.1111/ectj.12097.

[15]. Ribeiro, M.T., S. Singh, and C. Guestrin, *"Why should I trust you?": Explaining the predictions of any classifier*, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2016: San Francisco, CA. p. 1135-1144. DOI: https://doi.org/10.1145/2939672.2939778.

[16]. Lundberg, S.M. and S.-I. Lee, *A unified approach to interpreting model predictions*, in *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. 2017. Vol. 30, p. 4765-4774.

[17]. Kapoor, S. and A. Narayanan, *Leakage and the reproducibility crisis in ML-based science*. Patterns, 2023. 4(9): p. 100804. DOI: https://doi.org/10.1016/j.patter.2023.100804.

[18]. Babii, A., E. Ghysels, and J. Striaukas, *Econometrics of machine learning methods in economic forecasting*, in *Handbook of research methods and applications in macroeconomic forecasting*, M.P. Clements and A.B. Galvão, Editors. 2024, Edward Elgar: Cheltenham. p. 246-273. DOI: https://doi.org/10.4337/9781035310050.00014.

[19]. Farmer, J.D. and D. Foley, *The economy needs agent-based modelling*. Nature, 2009. 460(7256): p. 685-686. DOI: https://doi.org/10.1038/460685a.

[20]. Giannone, D., L. Reichlin, and D. Small, *Nowcasting: The real-time informational content of macroeconomic data*. Journal of Monetary Economics, 2008. 55(4): p. 665-676. DOI: https://doi.org/10.1016/j.jmoneco.2008.05.010.

[21]. Lazer, D., R. Kennedy, G. King, and A. Vespignani, *The parable of Google Flu: Traps in big data analysis*. Science, 2014. 343(6176): p. 1203-1205. DOI: https://doi.org/10.1126/science.1248506.

[22]. Angrist, J.D., G.W. Imbens, and D.B. Rubin, *Identification of causal effects using instrumental variables*. Journal of the American Statistical Association, 1996. 91(434): p. 444-455. DOI: https://doi.org/10.2307/2291629.

[23]. Card, D., *The impact of the Mariel Boatlift on the Miami labor market*. Industrial and Labor Relations Review, 1990. 43(2): p. 245-257. DOI: https://doi.org/10.2307/2523702.

[24]. Card, D. and A.B. Krueger, *Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania*. American Economic Review, 1994. 84(4): p. 772-793. Available from: http://www.jstor.org/stable/2118030.

[25]. Pearl, J., *Causality: Models, reasoning, and inference*. 2nd ed. 2009, Cambridge: Cambridge University Press.

[26]. Imbens, G.W. and D.B. Rubin, *Causal inference for statistics, social, and biomedical sciences: An introduction*. 2015, Cambridge: Cambridge University Press. DOI: https://doi.org/10.1017/CBO9781139025751.

[27]. Lucas, R.E., Jr., *Econometric policy evaluation: A critique*. Carnegie-Rochester Conference Series on Public Policy, 1976. 1: p. 19-46. DOI: https://doi.org/10.1016/S0167-2231(76)80003-6.

[28]. Akerlof, G.A., *The market for "lemons": Quality uncertainty and the market mechanism*. Quarterly Journal of Economics, 1970. 84(3): p. 488-500. DOI: https://doi.org/10.2307/1879431.

[29]. Stiglitz, J.E., *Information and the change in the paradigm in economics*. American Economic Review, 2002. 92(3): p. 460-501. DOI: https://doi.org/10.1257/00028280260136363.