

From Prediction to Causality in Data Science

1.1 Introduction**1.1.1 Evolution of Data Science: From Description to Prescription**

There are three significant phases of data science development. The first phase was on descriptive analytics where statistical software were employed to summarize past data and report what has already occurred. This involves standard examination of economic indicators like rates of growth, inflation and employment rates in economics. The second phase presented predictive analytics, which was motivated by progress in machine learning and massive data availability. In this case, models are conditioned to predict the future, which is market tendencies, demand, or macroeconomical variables.

In more recent times the emphasis has moved on to prescriptive analytics besides describing and predicting outcomes it tries to aid in decision-making. Prescriptive methods attempt to prescribe action by integrating evidence-based models and decision goals. This shift is indicative of a wider trend that data science must not only analyze the world, but also assist in transforming it in a favorable manner. Here, the role of data science is not only growing out of a passive analysis tool but also an active part of a decision-making system, particularly in the economic and public policy world (Lo and Pachamanova, 2023) [1]. Figure 1.1 is an image that represents the evolution of data science paradigms between descriptive and causal data science.



Figure 1.1: Evolution of Data Science

1.1.2 Limitations of Correlation-Based Prediction

Although predictive models have been successful, the majority of the existing methods are still constructed on statistical relationships and not cause-effect associations. Models trained on correlation can work well with the data they are trained in, however they exhibit a strong tendency to fail when applied to direct real-life interventions. The reason is that this type of models fails to differentiate between variables that move concurrently and variables that affect one another.

In most real-life scenarios, particularly in economics, one is required to make decisions about a system, but not observe it. As an illustration, the change in underlying data-generating process is the introduction of a new tax policy or a social welfare program. The result of such changes can not be accurately estimated using a model based on the historical correlations. Consequently, correlation-based prediction is prone to making false conclusions when used to assess policies or plan strategies (Cox Jr, 2024) [4].

The other drawback is that most predictive models that have been identified to be high-performing are black-box. Although they might be very good at numerical performance, they do not give much insight into the reasons as to why a specific outcome is expected. This is not very transparent and this reduces their utility in policy matters, where decisions need to be explained and with justifiable reasons and not obscure algorithmic results. Table 1.1 summarizes descriptive, predictive and causal analysis paradigms in a comparative way.

Table 1.1: Comparison of Different Data Analysis Paradigms

Aspect	Descriptive Analysis	Predictive Analysis	Causal Analysis
Main question	What happened?	What will happen?	What happens if we change X?
Primary goal	Summarize and understand past data	Forecast future outcomes	Estimate cause-effect relationships
Typical methods	Statistics, dashboards, reports	Machine learning, regression, deep learning	Causal graphs, structural models, counterfactual analysis
Output	Trends, patterns, summaries	Predictions, risk scores, probabilities	Causal effects, treatment impacts, policy outcomes
Role in decision-making	Informational	Support planning	Support interventions and policy design
Ability to	No	Very limited	Yes (explicitly)

answer “what if” questions			
Risk of misleading conclusions	Low (descriptive only)	Medium (correlation-based)	Lower (if assumptions hold)

1.1.3 Importance of Causality in Economics and Policy

Economics has been aware of the fact that sound policy analysis needs causal reasoning. The issue of many economic research is not merely what correlates to what, but what is the impact of an intervention. As an example, in evaluation of an education reform or a labour market policy, it is often desired to establish whether the policy in question influences improvements in outcomes, rather than to know whether the participants are different than the non-participants.

Causal models give a framework of distorting real effects and confounding effects. This is necessary in answering questions related to policies like whether a program should be increased, changed, or abandoned. Even advanced statistical models would not be able to provide credible guidance without causal identification to aid in decision-making. In this respect, prediction and explanation are utilized differently: prediction is concerned with accuracy, whereas causation is concerned with comprehension and manipulation (Heckman and Pinto, 2022) [2].

According to the latest studies, machine learning techniques can be applied to economic analysis more efficiently in the case they are united with the principles of causal inference. Rather than regarding models as the mere tools of prediction, this integration enables them to play a part in more credible estimation of the impact of policies and more justifiable economic findings (Ahrens et al., 2020) [5].

1.1.4 Toward Causal Data Science and Causal Economics

There has been an increase in the popularity of causality, and this has resulted in the development of causal data science, a discipline that seeks to bring together machine learning, statistics, and causal inference. This is in contrast to pattern recognition by focusing on mechanisms that produce data and the outcomes of interventions. Recent surveys indicate the use of causal reasoning that can enhance model robustness, feature selection, and interpretability and thus data-driven systems become more reliable in the real world (Lamsaf et al., 2025) [3].

Economists regard this change as favoring the emergence of causal economics, in which causal models become the foundation of policy-making and empirical research. The important rationale is just simple that economic decisions are actions in complex systems and that causal knowledge is the only one that can warrant those actions.

Causal models of outcomes are more appropriate in assessing trade-offs, predicting unintended outcomes, and in the implementation of effective policies. The movement is also in line with the larger issues of transparency, trust, and accountability in the use of data in decision-making. When conditions shift, systems which are based on correlations only can do well but fail unexpectedly. In comparison, causally informed models provide more convenient to interpret, communicate, and defend explanations in high-stakes economic and policy settings (Lamsaf et al., 2025) [3].

1.1.5 Motivation and Scope of This Chapter

Combining all these facts, the development of data science, the restrictions of prediction based on correlation, and the needs of economic policy make one obvious conclusion: causality should take centre stage in data-oriented decision-making in recent times. The transition between prediction and causation is not just a technical enhancement but a massive shift in the way models can be applied in order to provide knowledge and action.

The chapter also provides this metamorphosis by stating why causal reasoning is necessary in economics and policy, as well as why conventional predictive methods cannot stand alone. It establishes the conceptual base of the remainder of the book, which is going to delve into how causal techniques, counterfactual logic, and analysis of interventions may be incorporated into data science to create more dependable, interpretable, and influential economic models.

1.2 Prediction vs Explanation

1.2.1 Predictive Accuracy vs Causal Understanding

The machine learning systems today have been used with great success in predicting tasks in a wide variety of domains. Their advantage is that they recognize patterns in bulk data and apply the patterns to produce the correct forecasts. But causal knowledge is not always associated with high predictive accuracy. A model is able to make predictions with accuracy of an outcome without being aware of what variables actually make the outcome or how modifications in the policy or environment would change the outcome.

The difference is more crucial especially when decision-making is involved and not forecasting only. According to Pichler and Hartig (2023) [8], the majority of predictive models are not created to estimate causal effects, but to optimize the error in prediction. Consequently, they can do well in constant contexts but not when answering what if questions regarding interventions. Likewise, Baiardi and Naghi (2024) [10] demonstrate that, although machine learning is capable of value addition to empirical analysis, their causal inference has to be further structured and is not limited to the typical predictive modeling.

This distinction is critical in practice, in economics and policy analysis. Predictive models can tell us the most probable, but they cannot themselves, answer, why it occurs or what will be different with the introduction of a new policy. It is this disparity between prediction and explanation that encourages the increasing popularity of causal machine learning methods. Figure 1.2 depicts the conceptual difference of predictive and causal modeling.

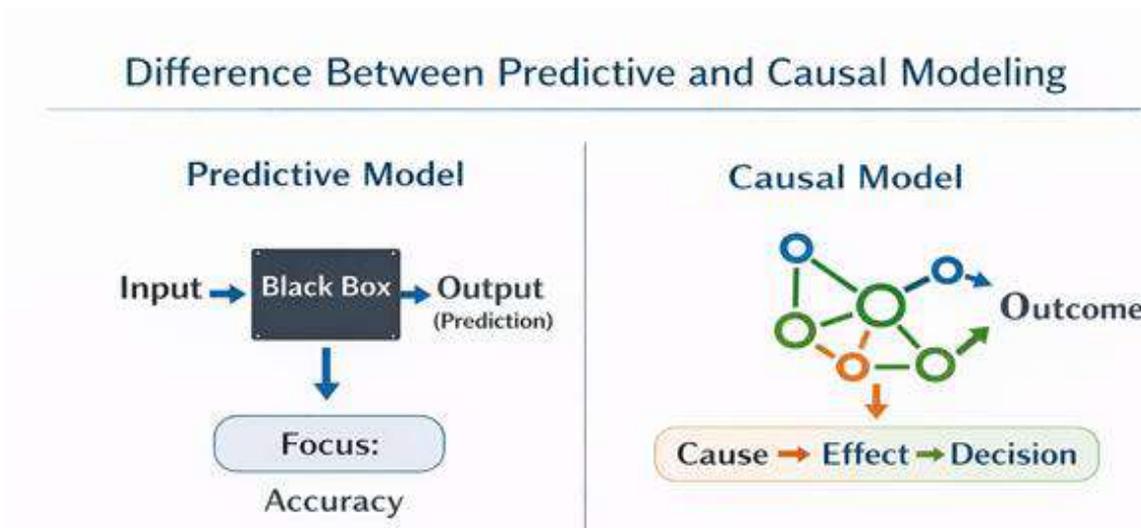


Figure 1.2: Prediction vs Causation

1.2.2 Black-Box Models vs Interpretable Causal Models

Most successful machine learning models such as deep neural networks and ensemble methods can be said to be black boxes since it is hard to tell how their internal decisions were reached. Although these models are capable of making valid predictions, their non-transparency hinders their usage in situations with a high stake where decisions should be justified and clarified.

However, causal models seek to capture explicit relationships between variables and thus, they allow one to trace how changes in one factor result in another. Renero et al. (2025) [9] suggest that machine learning can be used together with explanatory methods in order to facilitate the process of causal discovery and demonstrate that interpretability is a critical step in the evolution between pattern recognition and causal interpretation. On the same note, Wyrembek et al. (2025) [7] show that causal machine learning frameworks are able to do much more than risk prediction to include intervention planning, and the practical value of interpretable causal structures.

Horton (2024) [6] goes on to state that causal economic machine learning is a way to bridge the gap between human reasoning and artificial intelligence and concentrates on

actionable, explainable models. Such a move away of opaque predictors into interpretable causal models is necessary when models are applied to support economics and policy decisions, in which the understanding of the mechanism of an outcome is as valuable as the prediction of the outcome. The central dissimilarities between predictive and causal modeling methods are outlined in Table 1.2.

Table 2: Predictive Models vs Causal Models

Feature	Predictive Models	Causal Models
Main objective	Maximize prediction accuracy	Identify cause–effect relationships
Focus	Correlation and patterns in data	Mechanisms and interventions
Model type	Often black-box (e.g., deep learning)	Usually structured and interpretable
Interpretability	Often low	Usually higher
Use in policy and decision-making	Limited and risky	Highly suitable
Ability to answer “what if” questions	No (or unreliable)	Yes (core purpose)
Robustness to distribution changes	Often weak	Generally stronger if causal structure is correct
Risk of spurious conclusions	High	Lower (with correct assumptions)

1.2.3 Examples from Economics: Growth, Inflation, and Unemployment

The difference between prediction and explanation can be easily seen in the economic applications. As an example, the machine learning models can be trained to predict the GDP growth, inflation rates, or unemployment rates with high accuracy with the help of historical data. Nonetheless, these predictions do not necessarily show, to which policies or economic aspects these are ascribed.

Policy wise, the most important questions are causal, what is the impact of an interest rate change on inflation? The question is how does public investment impact economic growth? Will labor market reform help decrease unemployment? These questions cannot be answered using models that assess correlations only, but rather interventions have to be evaluated. The machine learning, as demonstrated by Baiardi and Naghi

(2024) [10], is capable of assisting causal analysis in economics, but it is possible only when incorporated into the causal inference model.

This is further confirmed by Wyrembek et al. (2025) [7] stating that causal machine learning allows predicting and planning because it explicitly models cause and effect, which are essential in the design and assessment of economic interventions. This way, causal approaches are not based on the ability to predict the economic indicators, but rather on explaining factors driving this and estimating the impacts of the policy actions.

To conclude, predictive models and causal models are used in different purposes. Predictive accuracy concentrates on the prediction of outcome whereas causal understanding assessment is used to explain mechanisms and to examine interventions. Black-box models can be excellent in prediction, but explanations, accountability and policy design require interpretable causal models. This difference is important in economics, where decisions are meant to alter variables like growth, inflation, and employment, and so on, it is obvious that good prediction does not imply good explanation, which supports the necessity of causal data science and causal economic modeling.

1.3 Failures of Pure Prediction

It is observed that purely predictive approaches might show statistical association, but they do not intrinsically test whether such association is actually associated with any causal mechanism. Even though it is good to have high accuracy in prediction, it is important to understand that it is not automatically good for understanding and decision-making in complex social, economic, and biomedical sciences.

- **Spurious Correlations**

One huge failure of prediction-driven analysis is its susceptibility to spurious correlation. Ghose et al. (2024) [11] explain that spurious associations arise because hidden confounders, effects of data aggregation, model overfitting, and accidental co-movement in high-dimensional data result in spurious associations. Predictive models may pick up on these regularities and achieve very robust performance metrics, but the regularities learned by the model may not reflect any actual causal process. Consequently, decisions based upon such models can turn out to be unstable once the conditions change or when models are applied outside their original context. This shows that correlation-driven prediction, even when strong statistically, gives a false sense of understanding and may lead to untrustworthy conclusions in practical applications.

- **Simpson’s Paradox**

Yet another well-known limitation of purely predictive forms of reasoning is demonstrated in Simpson’s paradox, which is observed in situations where, after aggregated trends in the given data are observed, the data is split into relevant categories, and the trends switch or vanish. The authors in Dong et al. (2024) have demonstrated that Simpson’s paradox cannot be caused by confounding and therefore demonstrates the inherently complicated nature of the relationships between given variables, with a predictive model formed on aggregated data potentially learning incorrect relationships, or incorrect and contradicting relationships, between the variables in question.

- **Policy Failures Due to Non-Causal Predictions**

Pure prediction has the most apparent limitations in the context of policy and decision making. According to Busetti (2023) [13], policies formulated based on predictive relations only expose themselves to failure since they fail to consider the causal processes that produce realized results. When interventions are based on correlation, there is a high probability that no improvements or even negative results will be generated. Similarly, Zawadzki et al. (2023) [14] mention that predictive associations, as applied in observational settings, cannot inform the action and causal models, such as confounder adjustment models or instrumental variable models are necessitated to estimate the actual action of interventions. These articles collectively underscore the idea that although predictive models can be used to inform the most likely occurrence, they fail to answer what ought to be done to justify the importance of causal reasoning in high-stakes decision-making. Some typical instances of policy risks that can occur due to non-causal, prediction-oriented reasoning are summarized in Table 1.3.

Table 1.3: Examples of Policy Risks When Using Non-Causal Predictions

Area	What was predicted	What was assumed (correlation)	What went wrong (no causality)	Lesson learned
Public health	Higher hospital visits predicted from certain behaviors	Behavior directly causes health outcome	Hidden confounders influenced both	Correlation is not causation
Economics	GDP growth predicted from selected indicators	Indicators were treated as causes	Policy based on indicators failed	Need causal validation before policy

Education	Student performance predicted from school features	Features assumed to cause performance	Socio-economic factors were ignored	Interventions must target true causes
Social policy	Crime predicted from demographic variables	Demographics assumed causal	Biased and ineffective interventions	Causal analysis is necessary for fairness
Business policy	Sales predicted from advertising patterns	Advertising assumed always effective	Market trends were the real driver	Decisions need causal impact estimates

1.4 Emergence of Causal Data Science

1.4.1 Judea Pearl’s Ladder of Causation

The growing disadvantages of pure predictive methods have prompted a revived interest in the significance of causality in data science. A method of interpreting the shift to a causal mode of reasoning is to refer to the three levels of reasoning that Judea Pearl has developed in her ladder of causation, which include association, intervention, and counterfactuals. The first level of reasoning involves the identification of associations, which allows a data science system to answer the question of what is likely to happen. Such reasoning is consistent with most of the conventional machine learning systems. Figure 1.3 is the summary of the three stages of reasoning as shown by the ladder.



Figure 1.3: Pearl’s Ladder of Causation

The upper stages of the ladder are concerned with intervention and counterfactual thinking. The intervention-based reasoning answers questions concerning the impacts of actions, including policy alterations or treatment whereas counterfactual reasoning evaluates hypothetical alternatives and questions what might have happened under varied circumstances. To ascend this ladder, it is needed that a causal structure is modeled and not merely statistical regularities. This trend is also evident in recent advances in causal machine learning, with the authors becoming more and more open to integrating structural causal concepts with data-driven approaches to justify their decision-making (Abécassis et al., 2025 [15]; Yu and Xu, 2025 [17]).

This theoretical construct has also had an effect on the creation of deep causal learning, in which representation learning, causal discovery and causal inference are addressed as systems. As demonstrated by Deng et al. (2025) [18], these methods are meant to go beyond the surface-level correlation and to the models that can embody both the generative and causal processes. These changes suggest that the contemporary data science is slowly ascending the causation ladder, and it is replacing the association with intervention and explanation.

1.4.2 From “What Will Happen?” to “Why Did It Happen?”

With this conceptual change, the objectives of data science are evolving as well. Classical predictive analytics mainly tries to respond to the question What will happen? by extending past trends with the help of historical data. Although this is good to predict, it gives minimal information when one wants to create interventions or analyze policies. By contrast, causal data science tries to address the question of Why did it happen? and What would happen if we changed something? and puts explanation and action in the focus of the analysis. According to Abercassis et al. (2025) [15] the transition involves a shift towards prescription in which machine learning models are being progressively applied to approximate the consequences of choices instead of just predicting them.

On the same note, Yu and Xu (2025) [17] illustrate that a combination of structural causal and machine learning facilitates the estimation of heterogeneous treatment effects, which is a direct support in making a decision in complex environments. These strategies represent how causal thinking can make data science into more than a descriptive and predictive instrument, and a model of how to comprehend processes and influence interventions.

It is also evidenced by new studies on causal foundation models. According to Zhang et al. (2023) [16], causal inference principles can enhance large-scale models by enhancing robustness, interpolability and generalization. This and additional progress in deep causal learning (Deng et al., 2025) [18] indicate that the future of the field of data science will not only be making predictions that are more accurate, but also

models capable of justification of the result and the interpretation of the implications of actions.

1.4.3 Causal Models as the New Foundation of Data Science

There is also the emergence of causal data science in the provision of new modeling paradigms that incorporate causal structure in the learning systems. According to Deng et al. (2025) [18], a holistic perspective of deep causal learning, representation learning, causal discovery, and causal inference is provided under cohesive frameworks. Their work demonstrates that the concept of causality is now being considered as a fundamental design concept, and is no longer being considered as a post-hoc interpretation added to black-box models. Simultaneously, causal foundation models research indicates a future where large-scale AI systems are capable of causal reasoning.

Zhang et al. (2023) [16] propose that the two are causal inference and attention processes, and suggest that future foundation models should be structured in such a way that causal reasoning is formalized in these models. This research supports the notion that effective generalization, interpretability and accurate decision-making entails more than precise prediction, but models capable of modeling and reasoning over cause-effect relationships. When combined, these trends imply that causality is coming to define modern data science. The discipline is shifting to correlations of a surface nature into models that can explain, perform intervention assessment, and reason about alternative possibilities, shifting the difference between predictive analytics and causal cognition.

1.5 Chapter Summary & Transition

1.5.1 Key Takeaways

The main thesis of this chapter has determined that although predictive data science is useful, it is not to be used in explanation, decision-making, or policy design. The contemporary data-driven systems are highly predictive, particularly with the development of AI and deep learning, but the predictions do not provide the mechanisms of how the results are created (Lo & Pachamano, 2023) [1]. Consequently, causal reasoning has to be used to complement prediction in the case where the aim is to justify interventions in the real world. One initial lesson learned is that prediction is helpful but not sufficient. Predictive models are useful to track trends and predict future conditions, but do not provide an answer to intervention-related questions, like what will happen in case a particular policy, treatment or decision is altered. In recent causal machine learning research, data science needs to be able to predict results, but it is required to help make decisions and actions as well (Abécassis et al., 2025 [15]; Wyrembek et al., 2025 [7]).

Second, the chapter has strengthened the idea that a correlation is not causation. Statistical associations can be seen as a result of confounding variables, selection effects, or structural biases and thus cannot be used as a foundation of policy or economic choices (Ghouse et al., 2024 [11]; Dong et al., 2024 [12]). However, without a clearly defined causal model, even very precise models may cause the wrong conclusion regarding the intervention effect. Third, economics and policy analysis are necessary because they involve causal thinking. Economic policies and choices are intervention-based: they attempt to assess the implications of actions instead of describing perceived trends. Causal econometric models and policy-practical causal study demonstrate that it is necessary to measure effects, not merely associations (Heckman and Pinto, 2022 [2]; Busetti, 2023 [13]).

Such a view is also justified by recent methodological developments that combine machine learning and causal inference to enhance decision relevance (Baiardi and Naghi, 2024 [10]; Ahrens et al., 2020 [5]). Lastly, it has been demonstrated in the chapter that data science is shifting towards causal reasoning as a principle. According to the recent surveys and methodological advances in causal learning, causal discovery, and deep causal models, a wider transformation of pattern recognition to explanation and intervention analysis is observed (Lamsaf et al., 2025 [3]; Deng et al., 2025 [18]; Yu and Xu, 2025 [17]). It is also this change that is closely associated with increased expectations of reliable and responsible AI systems whose justifications should be based on cause-effect relationships and not just correlations (Rawal et al., 2025 [20]).

1.5.2 Transition to Counterfactual Reasoning

Based on these findings, the next logical move is to go beyond general causal principles to the counterfactual reasoning. A counterfactual question is one that inquires: What could have been different had we done things differently? This line of thinking is the foundation of causal inference and modern policy assessment since it directly accounts for the impact of other decisions as opposed to mere forecasting of future consequences with all things else held constant (Shmueli et al., 2025) [19]. Whereas predictive models determine the probable results, the counterfactual analysis makes a comparison between hypothetical circumstances and assesses the effects of particular actions. It is this difference that leads to the shift of descriptive and predictive analytics to truly decision-oriented data science.

Recent literature emphasizes that counterfactuals are a key area of causal inference and explainable AI, albeit with different objectives: causal inference aims to estimate the effects of an intervention, and explainable AI may typically use counterfactuals to enhance interpretability and transparency (Shmueli et al., 2025) [19]. It is viewed as a foregone conclusion that, at least in high-stakes fields like economics and government policy, causal and counterfactual reasoning becomes required to achieve reliable,

accountable systems, based on the concept of trustworthy AI (Rawal et al., 2025) [20]. Hence, the chapter ends on the note that it prepares the ground on the next chapter that will dwell on the counterfactuals, interventions, and estimation of causal effects that will shift the conceptual arguments to tools and methods of making a causal decision.