

Article

An Institutional Theory Framework for Leveraging Large Language Models for Policy Analysis and Intervention Design

J. de Curtò ^{1,2,3,*} , I. de Zarzà ^{3,4} , Leandro Sebastián Fervier ⁵ , Victoria Sanagustín-Fons ⁵ 
and Carlos T. Calafate ⁶ 

- ¹ Department of Computer Applications in Science & Engineering, BARCELONA Supercomputing Center, 08034 Barcelona, Spain
 - ² Escuela Técnica Superior de Ingeniería (ICAI), Universidad Pontificia Comillas, 28015 Madrid, Spain
 - ³ Estudis d'Informàtica, Multimèdia i Telecomunicació, Universitat Oberta de Catalunya, 08018 Barcelona, Spain; izarza@unizar.es
 - ⁴ Departamento de Informática e Ingeniería de Sistemas, Universidad de Zaragoza, 50009 Zaragoza, Spain
 - ⁵ Departamento de Psicología y Sociología, Universidad de Zaragoza, 50009 Zaragoza, Spain; 848408@unizar.es (L.S.F.); vitico@unizar.es (V.S.-F.)
 - ⁶ Departamento de Informática de Sistemas y Computadores, Universitat Politècnica de València, 46022 València, Spain; calafate@disca.upv.es
- * Correspondence: jdecorto@icai.comillas.edu

Abstract: This study proposes a comprehensive framework for integrating data-driven approaches into policy analysis and intervention strategies. The methodology is structured around five critical components: data collection, historical analysis, policy impact assessment, predictive modeling, and intervention design. Leveraging data-driven approaches capabilities, the line of work enables advanced multilingual data processing, advanced statistics in population trends, evaluation of policy outcomes, and the development of evidence-based interventions. A key focus is on the theoretical integration of social order mechanisms, including communication modes as institutional structures, token optimization as an efficiency mechanism, and institutional memory adaptation. A mixed methods approach was used that included sophisticated visualization techniques and use cases in the hospitality sector, in global food security, and in educational development. The framework demonstrates its capacity to inform government and industry policies by leveraging statistics, visualization, and AI-driven decision support. We introduce the concept of “institutional intelligence”—the synergistic integration of human expertise, AI capabilities, and institutional theory—to create adaptive yet stable policy-making systems. This research highlights the transformative potential of data-driven approaches combined with large language models in supporting sustainable and inclusive policy-making processes.

Keywords: data-driven policy analysis; institutional theory; decision support systems; visual analytics; AI; large language models; predictive modeling; intervention design; graph neural networks



Academic Editors: Filipe Portela and Athanasios D. Panagopoulos

Received: 6 January 2025

Revised: 13 February 2025

Accepted: 17 February 2025

Published: 20 February 2025

Citation: de Curtò, J.; de Zarzà, I.; Fervier, L.S.; Sanagustín-Fons, V.; Calafate, C.T. An Institutional Theory Framework for Leveraging Large Language Models for Policy Analysis and Intervention Design. *Future Internet* **2025**, *17*, 96. <https://doi.org/10.3390/fi17030096>

Copyright: © 2025 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The rapid advancements in artificial intelligence (AI) have led to transformative shifts in how societies interact, organize, and manage information. Among these advancements, large language models (LLMs) [1] are transformative tools capable of processing vast amounts of data (e.g., legal corpora, government data), identifying patterns, and generating actionable insights. This study introduces a novel framework that leverages the capabilities

of advanced data-driven approaches to address critical aspects of policy analysis and intervention design [2]. The framework is structured around five interrelated components: data collection, historical analysis, policy impact assessment, predictive modeling, and intervention design, each of which is underpinned by the ability of data-driven approaches to adapt and optimize across diverse contexts [3].

Communication modes within data-driven systems are conceptualized as institutional structures that play a fundamental role in maintaining social order. These modes enable the efficient organization and interpretation of information while preserving the quality and integrity of interactions. By exploring efficiency mechanisms such as token optimization and information density, the framework ensures that meaning and context are preserved even across diverse communication types. This approach not only enhances the responsiveness of data-driven systems, but also fosters trust and legitimacy in their outputs [4].

Token optimization, in this context, refers to a computational mechanism to streamline information processing within large language models and other AI-driven analytical tools. By efficiently managing the “tokens”—the basic units of text or data the model consumes—we enhance processing speed, reduce computational overhead, and preserve interpretive quality. Briefly, token optimization ensures that the model maintains clarity and relevance in its outputs while minimizing redundant or extraneous information. Thus, this efficiency mechanism supports institutional consistency and trustworthiness in AI-driven policy analysis, ensuring that linguistic complexity does not undermine interpretative coherence or data integrity.

For example, imagine a regional environmental agency deciding how to allocate limited water resources among agricultural, industrial, and residential users. By employing token optimization and structured communication protocols, an AI-driven decision support system can efficiently process large volumes of multilingual policy documents and stakeholder testimonies, then generate concise, context-aware recommendations. This ensures that rather than sifting through dense policy texts or facing ambiguous guidance, decision-makers receive clear, prioritized policy options that respect institutional rules, improve transparency, and ultimately enhance public trust in the allocation process.

This research seeks to address three central questions:

1. How do data-driven entities’ communication modes function as institutional structures to maintain social order?
2. What efficiency mechanisms enable these systems to optimize responses without compromising quality?
3. How do these adaptive systems reshape communication hierarchies and institutional memory in policy-making?

Drawing on institutional theory and social systems analysis, this study provides a multidisciplinary perspective on integrating data-driven approaches into frameworks for societal decision-making. By examining theoretical implications, including the impact on communication hierarchies and adaptive responses, this framework offers insights into how systems governed by data can be leveraged to design and implement evidence-based interventions while maintaining social cohesion.

Public policies consist of a set of rules that determine, limit, incentivize, and influence decision-making by balancing costs and benefits in the allocation of political and economic resources, ultimately shaping societal development [5]. To analyze public policies, a theoretical framework is essential. This paper adopts the “process sequencing” approach [6,7], which posits that events are causally connected within a public policy sequence [8–11].

It can be argued that governments choose actions aimed at addressing issues identified as central, enabling analysts to perceive the phenomenon that requires explanation. The

rational choice of an agent seeks to maximize values by selecting the intervention alternative that best facilitates the achievement of their goals and objectives [12]. However, it is essential to move away from the notion of absolute rationality in decision-making processes as public policy processes reveal, to some extent, how governments pursue certain objectives but not why those objectives are prioritized [13]. In summary, decisions are interconnected operations that involve selecting an alternative from a set of options, or more simply, the choice of a course of action [14].

In the context of policy analysis, institutional memory [15] refers to the accumulated knowledge, practices, and decisions that shape an organization's long-term strategy. It includes formal elements such as archives, records, and databases, as well as informal aspects like shared experiences, cultural norms, and expert intuition within policy institutions. The analogy between decision-making processes [16] can lead to repeating the formula that has been successfully applied in the past. Our framework operationalizes institutional memory through historical analysis and predictive modeling, ensuring that past policy successes and failures inform present decision-making. Organizational adaptability [17,18] refers to an institution's ability to modify strategies and decision-making structures in response to external changes, policy shocks, and technological advancements. In our framework, adaptability is captured through predictive modeling, AI-assisted learning mechanisms, and intervention design, allowing organizations to respond proactively to emerging policy challenges while maintaining institutional stability.

To provide a clear roadmap, this paper is structured as follows. In Section 2, we review the foundational literature in institutional theory and its applications to social systems, establishing the theoretical groundwork for our framework. Section 3 introduces our comprehensive framework for data-driven policy analysis, detailing its five core components and communication structure. In Section 4, we bridge theoretical foundations with practical applications, demonstrating how our methodology translates into concrete tools and processes. Section 5 showcases practical implementations in the hospitality sector, featuring advanced visualization techniques, graph neural networks, and LLMs. Section 6 further explores additional use cases from related policy domains: global food security and educational development. We then present our discussion in Section 7, analyzing the framework's effectiveness and implications across multiple dimensions. Finally, in Section 8, we conclude by summarizing our contributions and highlighting the framework's potential for broader application across different policy domains. Each section builds upon the previous ones to demonstrate how data-driven approaches can be effectively integrated with institutional theory to create robust, actionable policy insights while maintaining social order cohesion.

2. Related Work

The foundations of institutional theory and its application to social systems are well documented in the literature. Elinor Ostrom's seminal work on the governance of commons [19] provides a detailed examination of how collective action and institutional arrangements influence resource management. Her framework for understanding the evolution of institutional rules and their enforcement mechanisms is central to much of the current discourse on institutional adaptation in policy contexts.

The exploration of partial organization, as discussed by Ahrne and Brunsson [20], extends these principles by examining how non-traditional organizational forms operate outside conventional boundaries. This concept aligns with the adaptive capacities of data-driven systems, which exhibit partial organization through decentralized, rule-based interactions, allowing for dynamic decision-making in complex environments.

Niklas Luhmann's work on social systems theory [21] offers a foundational perspective on the relational and processual nature of social order. Luhmann posits that systems maintain their integrity through communication, which resonates with the role of AI in sustaining institutional trust through consistent and high-quality outputs. These systems, as Luhmann theorizes, are defined by their ability to self-regulate and adapt, a characteristic mirrored in modern AI frameworks.

Building on these theoretical insights, Berkowitz and Grothe-Hammer [22] investigate the dynamics of social orders and the challenges posed by meta-organizations in global contexts. Their analysis of decidability in organizational frameworks highlights the critical role of power and decision-making structures in maintaining institutional legitimacy. This directly informs the design of AI systems that aim to facilitate equitable policy outcomes through adaptive communication modes.

Finally, Scartascini et al. [23] provide a political economy perspective on decision-making in Latin America, emphasizing the interplay of institutions, networks, and strategic interactions. Their findings note the importance of contextual sensitivity in policy design, a feature that visualizations, AI, and statistics can enhance through their ability to process and synthesize diverse inputs across multilingual and multi-sectoral datasets.

Recent advancements in AI for decision-making have demonstrated the value of hybrid frameworks integrating visualization, statistical inference, and machine learning. Decision support systems (DSSs) [24,25] and evidence-based policy (EBP) [26–28] have traditionally been employed in policy-making, but recent research has explored the integration of AI for predictive analytics and enhanced policy recommendations [29,30]. In particular, LLMs have been shown to improve policy formulation through natural language synthesis and structured guidance [31]. This study builds upon these works by proposing an integrated framework that combines AI, visualization, institutional theory, and predictive analytics to enhance institutional decision-making.

These works collectively establish a robust theoretical and practical foundation for integrating data-driven approaches into frameworks for policy analysis and intervention. By synthesizing these insights, the proposed framework seeks to operationalize the principles of institutional theory, efficiency mechanisms, and social order integration in AI-driven systems.

3. Materials and Methods

The proposed framework leverages data-driven approaches to address complex policy analysis and intervention challenges. The methodology is structured around five interconnected stages: data collection, historical analysis, policy impact assessment, predictive modeling, and intervention design. Each stage is designed to optimize the integration of visualization, statistics, and AI capabilities, ensuring efficiency and maintaining social order through adaptive communication and decision-making processes.

The methodological framework of this study is primarily based on the theoretical principles of Luhmann and Ostrom. For the technical implementation, we utilized several generative AI tools: ChatGPT/Copilot/Gemini/Claude were employed to assist with code completion and generate initial pseudocode structures, which were subsequently reviewed, modified, and validated by the authors. These tools were specifically used to optimize the diagramming processes and architect the interactive components, as well as for preprocessing of the data needed for the visualizations. All generated code and diagrams underwent thorough human review and testing to ensure accuracy and functionality.

In social sciences, a “case study” traditionally refers to a detailed, often qualitative examination of a single instance or setting. In contrast, our usage of “use cases” aligns more closely with an applied problem-solving perspective drawn from fields such as information

systems, computer science, and organizational management. By “use cases”, we mean specific illustrative examples or scenarios that demonstrate how our framework’s components and analytical tools (visualizations, models, and interventions) can be operationalized in practice. Unlike the in-depth focus of a single case study, multiple “use cases” here reflect a series of targeted, structured examples, each highlighting different functionalities of the framework, ranging from historical data analysis to predictive modeling and policy intervention design. We retain the term “case study” where appropriate to align with established qualitative traditions, while “use cases” signal more modular, application-oriented vignettes that complement our overarching methodology.

3.1. Framework Overview

Figure 1 presents a high-level overview of the framework, illustrating the sequential progression from data collection to intervention design. This figure provides a clear, simplified representation of the primary stages involved in our approach.




Figure 1. High-level framework for policy analysis and intervention design. This diagram shows the core progression through the five main stages of the framework.

Figure 2 expands on this structure by detailing the specific components and interconnections within each stage. This enhancement facilitates a more granular understanding of how data flow through the framework, highlighting the relationships between different analytical processes and decision-making mechanisms.




Figure 2. Detailed components of the framework stages. This diagram details the specific elements and interconnections within each framework stage.

The data collection phase synthesizes multilingual data sources, including demographic data, policy documents, and local testimonies. These inputs are processed using statistics and AI to extract structured information, ensuring a comprehensive dataset for analysis.

Using the collected data, historical trends such as population changes, policy timelines, and economic factors are analyzed to identify critical patterns. This phase informs the policy impact assessment by providing context for short- and long-term outcomes.

This stage evaluates the potential and realized effects of policies. Metrics such as short-term effects, long-term outcomes, and success metrics are analyzed to quantify the impact and provide comparative insights.

Predictive modeling employs visualization techniques, statistics, and AI to forecast population trends, assess risks, and identify resource needs. The insights generated in this phase serve as a foundation for designing evidence-based interventions.

The final phase focuses on translating insights into actionable plans. Policy recommendations, implementation strategies, and monitoring systems are developed to ensure effective interventions.

3.2. Communication Structure

To maintain social order, the framework incorporates a robust communication structure. Figure 3 presents the high-level architecture, showing how three core layers—determination, efficiency, and interaction—contribute to output formation and ultimately impact social order. This figure captures the essential flow of information and decision-making within the framework.




Figure 3. High-level communication structure. This diagram presents an overview of how communication structure integrates determination, efficiency, and interaction layers, leading to output formation and social order impact.

Figure 4 details the specific components within each layer, illustrating their hierarchical relationships and interconnections. The upper section of this figure focuses on the determination layer, which includes rule-based responses, contextual adaptation, and mode selection, providing structured and contextually relevant outputs. The middle section details the efficiency layer, emphasizing token optimization, information density, and quality preservation to ensure meaningful and resource-efficient communication. The lower section highlights the interaction layer, incorporating user feedback loops, adaptive responses, and mode transparency to enhance trust and legitimacy in AI-driven interventions.




Figure 4. Detailed layer components of communication structure. This diagram expands on the three primary layers—determination, efficiency, and interaction—detailing their internal components and interconnections.

The determination layer includes rule-based responses, contextual adaptation, and mode selection, enabling visualization techniques, statistics, and AI to provide structured and contextually relevant outputs.

This layer focuses on token optimization, information density, and quality preservation, ensuring that outputs are both resource-efficient and meaningful.

The interaction layer incorporates user feedback loops, adaptive responses, and mode transparency to enhance trust and legitimacy in AI-driven interventions.

3.3. Validation Procedures

The framework's validity can be assessed through:

- Validation of metrics across multiple datasets and scenarios.
- Expert review of qualitative outputs and framework applications.
- Triangulation of insights derived from visualization techniques, statistics, and AI outputs to ensure consistency and reliability.

Our mixed-methods approach integrates qualitative insights and quantitative metrics within a cohesive analytical framework. Qualitative elements include an interpretive modeling of institutional theory, examination of communication structures, and the use of socio-organizational concepts (e.g., Ostrom's collective action principles). Quantitative components encompass statistical trend analyses, advanced visualization (e.g., sunburst charts, geographic maps), and AI-driven modeling techniques like graph neural networks. By combining these methods, we leverage the strengths of both rigorous data-driven computation and rich, context-sensitive theoretical understanding. This synthesis allows us not only to identify patterns and predict outcomes but also to interpret these findings through the lens of institutional theory, thereby yielding more robust and actionable policy insights.

Our framework validation strategy emphasizes practical implementation in real-world scenarios, with a specific focus on the hospitality sector [32,33] as a representative case study. This domain was chosen for its complex interplay of temporal, spatial, and hierarchical factors—characteristics that align well with our framework's multi-layered approach. Through this implementation, we demonstrate how our theoretical components manifest in practical applications, from data collection through to intervention design, while incorporating advanced techniques such as graph neural networks into the dashboard, and we discuss the use of large language models.

3.4. Comparison with Existing Policy Analysis Frameworks

To highlight the advantages of our proposed framework, we compare it with existing policy analysis methodologies, particularly decision support systems (DSSs), evidence-based policy (EBP), and institutional modeling frameworks. Table 1 presents a structured comparison of these approaches.

Our framework extends traditional DSS models by incorporating LLM-powered analysis [31] to generate structured recommendations, ensuring contextual adaptation. Additionally, while EBP frameworks primarily rely on retrospective data, our approach integrates predictive analytics, enabling policy-makers to simulate interventions before implementation. This combination allows for an advanced decision-making paradigm that balances institutional stability with adaptive policy innovation.

The integration of AI, visualization techniques, and statistical modeling follows a structured analytical pipeline. Visualization enables pattern discovery and human interpretability, statistics provides rigorous empirical validation, and AI enhances predictive modeling and scenario-based decision-making. This multi-layered approach ensures that decision-makers receive actionable insights while maintaining institutional transparency and robustness. For example, in the hospitality use case, interactive visualizations identify pricing patterns, while graph neural networks model complex booking dynamics, and large language models synthesize policy recommendations grounded in institutional memory.

Table 1. Comparison of policy analysis frameworks.

Feature	Decision Support Systems (DSS) [24,25]	Evidence-Based Policy (EBP) [26–28]	Institutional Intelligence (Ours)
Integration of AI and LLMs	Partial	No	Yes (AI-driven modeling, LLM-powered recommendations)
Historical Memory Preservation	Low (data archiving only)	Medium (case studies)	High (institutional memory through structured AI inference)
Predictive Capabilities	Medium (basic forecasting)	Low (focus on retrospective data)	High (machine learning and graph neural networks)
Adaptability to New Policies	Medium	Low	High (AI-assisted scenario analysis, policy simulation)
Visualization and Explainability	Low	Medium	High (interactive visualizations, LLM-generated narratives)

The framework operates on structured time intervals based on the policy domain. For example, in the hospitality sector use case, booking data includes arrival dates (check-in), departure dates (check-out), and cancellation dates to capture customer behavior. Similarly, in the food security and education sectors, data are analyzed annually, focusing on trends across pre-crisis, crisis, and post-crisis periods (e.g., 2019–2022 for food security and 2000–2020 for education). These structured temporal dimensions allow for robust historical analysis, predictive modeling, and intervention design, ensuring that insights remain policy-relevant.

While this paper incorporates existing analytical techniques such as statistical modeling, visualization methods, and AI-based decision support, it proposes a novel framework that integrates institutional theory with data-driven approaches. Unlike prior studies that focus on either policy modeling or AI-driven decision support in isolation, our framework introduces the concept of institutional intelligence—a synergistic combination of human expertise, AI-assisted modeling, and structured policy analysis. This framework is not merely descriptive; it is prescriptive, providing policy-makers with a structured methodology for adapting AI-powered decision tools to institutional constraints, hierarchical decision-making, and policy design. To further emphasize this contribution, we introduce a step-by-step implementation guide outlining how each component of the framework can be practically deployed across various policy domains.

4. Framework Implementation and Validation Strategy

Before presenting specific use cases, it is essential to establish how the methodological framework translates into practical applications. This section bridges the theoretical foundation presented in Materials and Methods with the concrete implementations demonstrated in the use cases.

4.1. Integration of Framework Components

The framework’s five core components align with the visualization implementations in the following ways:

- Data collection: The hospitality sector implementation draws from multiple data sources, including booking records, pricing information, and geographical data, demonstrating the framework’s capacity for multi-source integration.

- Historical analysis: The temporal visualization techniques, particularly the violin plots and cumulative heatmaps, directly implement the framework's historical analysis component by revealing patterns across different time scales.
- Policy impact assessment: The interactive visualizations, especially the geographic mapping and sunburst diagrams, enable stakeholders to assess the impact of various market conditions and policy decisions on key metrics such as average daily rate (ADR) and cancellation rates. ADR is a key hospitality industry metric that represents the average revenue earned per available room per day.
- Predictive modeling: The temporal pattern analysis and clustering visualizations support the framework's predictive capabilities by highlighting recurring patterns and trends.
- Intervention design: The multi-layered visualization approach facilitates evidence-based intervention design by providing comprehensive views of market behavior and performance metrics.

4.2. Communication Structure Implementation

The three-layer communication structure described in the methodology manifests in the visualization strategy through:

1. Determination layer implementation
 - Rule-based responses are implemented through consistent visualization formats across different metrics.
 - Contextual adaptation is achieved through interactive filtering and drill-down capabilities.
 - Mode selection is demonstrated through multiple visualization types for different aspects of the data.
2. Efficiency layer implementation
 - Information density is optimized through the use of multi-dimensional visualizations.
 - Quality preservation is maintained through careful selection of visualization techniques.
 - Token optimization is achieved through intuitive visual hierarchies.
3. Interaction layer implementation
 - User feedback loops are enabled through interactive features.
 - Adaptive responses are implemented through dynamic filtering options.
 - Mode transparency is maintained through clear visual hierarchies and consistent design patterns.

4.3. Validation Approach

The framework's effectiveness is validated through:

- Multi-dimensional analysis: the use cases demonstrate the framework's ability to handle spatial, temporal, and hierarchical dimensions simultaneously.
- Scalability testing: the implementation across different metrics and time periods validates the framework's scalability.
- User interaction: the interactive features demonstrate the framework's capacity for dynamic analysis and adaptation.
- Information synthesis: the visualization stories show how complex data can be effectively communicated through structured narratives.

The subsequent sections present a detailed examination of this implementation strategy through specific use cases in the hospitality sector. These cases demonstrate not only the framework's practical utility, but also its alignment with the theoretical foundations

established earlier. By progressing from visualization techniques through advanced machine learning approaches like graph neural networks (GNNs)—that is, machine learning models designed for structured data, enabling advanced relational learning across interconnected policy datasets—and LLMs, we show how each theoretical component manifests in practical applications while maintaining the framework’s core emphasis on social order integration and institutional structure preservation.

In the next sections, we propose several use cases to strengthen the empirical validation of our framework, we incorporate a structured analysis aligned with the framework implementation and validation strategy, as well as advanced tailor-made interactive visualizations with sophisticated preprocessing of data from public datasets and LLM policy generation with LLMs. Specifically, for instance, we expand the hospitality sector case study by performing an ADR market segmentation analysis using graph neural networks (GNNs). This implementation aligns with our predictive modeling and intervention design components. Our results demonstrate that seasonal pricing fluctuations, customer segmentation, and regional influences are key determinants of market stability. These inclusions provide a data-backed validation of our framework, reinforcing its applicability beyond theoretical discussions. Additionally, we explore additional use cases from related policy domains where institutional intelligence emerges from advanced visualization techniques, with sophisticated prior preprocessing of data to extract a meaningful narrative and interactive plots coded with plotly in Python and presented as a public web app in Streamlit.

5. Use Cases

The following use cases demonstrate the practical implementation of our framework’s theoretical components in the hospitality sector. These implementations specifically address the research questions outlined in the introduction: (1) how AI communication modes function as institutional structures, demonstrated through our visualization hierarchy; (2) efficiency mechanisms in system optimization, shown through our GNN implementation; and (3) the reshaping of communication hierarchies, evidenced by our multi-layer analytical approach. Each visualization and analysis technique maps directly to one or more components of our theoretical framework, providing concrete examples of how abstract concepts translate into practical tools for policy analysis and intervention design.

Effective data analysis, visualizations, and statistics are fundamental to understanding the complex interplay of temporal, hierarchical, and spatial factors in decision-making and resource allocation. This study integrates advanced visualization techniques with data discovery to uncover patterns and provide actionable insights into the hospitality market [33]. Below, we describe each key data visualization story and its role in supporting our mathematical and analytical framework.

For our visual analytics, we utilized Flourish (web version), ensuring we worked with the latest cloud-based version available as of December 2024. Similarly, for Tableau-based visualizations, we employed Tableau Public (web version). Our interactive dashboard implementations were developed using Streamlit (web app), running on Python, and our data preprocessing and analysis were conducted using Python with essential libraries, including pandas for data manipulation, numpy for numerical computations, and plotly for interactive visualizations.

5.1. Temporal Insights Through ADR Variability: Violin Plot

This visualization implements the framework’s temporal analysis component while demonstrating how institutional structures (as discussed in our theoretical framework) can be maintained through consistent visual communication modes. The violin plot visualizes the distribution of average daily rate (ADR) across different time periods (e.g., months,

seasons, or years) and by hotel type (city versus resort). This plot combines features of boxplots and kernel density plots, showing not only summary statistics (e.g., medians, interquartile ranges) but also the full distribution of the ADR.

The violin plot highlights variability in pricing strategies, which is essential for:

- Temporal optimization: identifying high-variance periods where dynamic pricing strategies could stabilize revenue.
- Seasonal trends: revealing peak and off-peak pricing patterns.
- Economic implications: supporting models of temporal resource allocation by quantifying demand elasticity.

5.2. Hierarchical Market Structure: Interactive Sunburst Chart

Building on Luhmann's [21] concept of system self-regulation and Ostrom's [19] institutional arrangements theory, the sunburst chart provides a hierarchical visualization of market structure. By segmenting data into nested layers such as countries, hotel types, and customer categories, this visualization embodies the partial organization principles discussed by Ahrne and Brunsson [20]. The interactive nature of this visualization allows users to drill down into specific segments to explore their contributions to cancellations and revenue, directly implementing the framework's institutional memory adaptation component.

This visualization offers insights into decision-making at multiple levels:

- Hierarchical decision analysis: quantifying the impact of higher-level decisions (e.g., country policies) on lower levels (e.g., hotel and customer type).
- Segmentation metrics: supporting entropy-based measures to quantify customer diversity.
- Game-theoretic implications: providing an intuitive view of NASH equilibria in pricing strategies across different market levels.

5.3. Spatial Patterns and ADR Clustering: Interactive Geographic Map

Following Scartascini et al.'s [23] emphasis on contextual sensitivity in policy design, this interactive map visualizes spatial distributions of ADR and cancellation rates across geographic regions. The implementation aligns with Berkowitz and Grothe-Hammer's [22] analysis of meta-organizational frameworks, particularly in how regional decision-making structures influence market outcomes.

The spatial map highlights regional disparities and clustering effects:

- Geographic insights: identifying regions with high ADR and low cancellation rates as benchmarks for stable markets.
- Spatial dependencies: visualizing spatial autocorrelation to support econometric models of cancellations.
- Policy recommendations: highlighting areas for targeted interventions based on regional market conditions.

5.4. Dynamic Trends in Market Leadership: Animated Bar Race

The animated bar race shows the evolution of top countries contributing to cancellations over time. Bars represent cancellation counts, with their lengths and ranks changing dynamically as the timeline progresses.

This visualization reveals:

- Temporal dynamics: highlighting shifts in market behavior over time.
- Competitor analysis: comparing country-level performance in managing cancellations.
- Longitudinal patterns: supporting predictive models by visualizing trends in a dynamic context.

5.5. Worldwide Temporal Patterns: Cumulative Data Heatmap

Drawing on the process sequencing approach [6,7], the cumulative heatmap aggregates data over multiple years, demonstrating how events are causally connected within policy sequences. This visualization particularly addresses the institutional memory aspects highlighted in our theoretical framework, showing how patterns persist and evolve across time periods.

The heatmap provides:

- Global comparisons: enabling cross-regional benchmarking of performance metrics.
- Temporal coherence: supporting Fourier and ARIMA models by identifying consistent periodicities.
- Market stability: highlighting regions with consistent patterns, ideal for risk-averse investments.

5.6. Integration of Visualizations into the Analytical Framework

The visualizations serve as critical tools to validate and communicate the findings of the mathematical models, implementing the institutional structures concept central to our theoretical framework. This integration demonstrates how Luhmann's [21] principles of system integrity through communication can be practically implemented:

- Temporal dimension: violin plots and heatmaps inform temporal resource allocation models.
- Hierarchical insights: sunburst charts align with entropy measures and game-theoretic predictions.
- Spatial dynamics: interactive maps and clustering analyses validate spatial econometric models.
- Dynamic trends: animated bar races contextualize changes in market leadership and behavior over time.

This study employs multiple visualization tools to enhance policy analysis after sophisticated data preprocessing in Python. Flourish is used for interactive storytelling, ideal for dynamic data exploration. Tableau enables multi-dimensional policy impact assessments through hierarchical visualizations. Streamlit allows for real-time dashboard generation through the Python library plotly, facilitating interactive decision-making. These tools complement predictive AI techniques by providing human-interpretable outputs.

Figure 5, developed in Flourish, is a story that presents three interconnected insights into the hospitality market: (1) an ADR interactive map illustrating the spatial distribution of average daily rates (ADRs) over time, enabling the identification of pricing clusters and trends across regions; (2) an ADR violin plot comparison, showcasing the distinct distributional pricing behaviors of city versus resort hotels; and (3) worldwide temporal hotel patterns, represented as cumulative monthly data highlighting trends in ADR, cancellations, and hotel performance.

Figure 6, developed in Tableau, comprises five distinct acts: (1) a bar chart comparing total bookings and cancellations by hotel type, highlighting disparities between city and resort hotels; (2) a heatmap visualizing cancellation rates across global regions, identifying geographic hotspots with high cancellation frequencies; (3) a temporal bar chart showing seasonal variations in bookings and cancellations, emphasizing peak and off-peak behaviors; (4) a stacked bar chart that visualizes customer type distribution across cancellation rates, offering insights into demographic impacts; and (5) a ranking chart displaying countries with the highest cancellation rates.




Figure 5. Interconnected insights into the hospitality market. These visualizations collectively offer a comprehensive view of the spatial, typological, and temporal dimensions of the hospitality industry. Interactive version available at: <https://public.flourish.studio/story/2733675/>, accessed on 1 January 2025.




Figure 6. Visualization story with five acts developed in Tableau. This multi-layered narrative explores temporal, regional, and demographic aspects of cancellations in the hospitality sector. Interactive version available at: <https://public.tableau.com/app/profile/decurto/viz/Tendenciasdereserveshotelsilescancellacions/Story1>, accessed on 1 January 2025.

Figure 7 is an interactive sunburst plot developed in D3.js that segments hotel reservation cancellations hierarchically by customer type, hotel type, and countries. The sunburst design allows for an intuitive exploration of the proportional contributions of various segments to the overall cancellation patterns. The visualization highlights key insights, such as the predominance of transient customers and city hotels in overall cancellation rates, providing actionable intelligence for segmentation strategies and targeted interventions.



Figure 7. Interactive sunburst plot developed in D3.js of hotel reservation cancellations. Interactive version available at: <https://decurto01.netlify.app/>, accessed on 1 January 2025.

Advanced data analytics serve as integral tools for analyzing and communicating the complex dynamics of the hospitality industry. Figure 5 presents the first story, which integrates spatial, typological, and temporal dimensions of ADR trends to offer a holistic view of market dynamics. Figure 6 unfolds a Tableau-based narrative in five acts, exploring diverse perspectives on cancellations and bookings, from demographic segmentation to geographic disparities and seasonal behaviors. Finally, Figure 7 introduces an interactive sunburst chart that enables hierarchical segmentation of cancellations, revealing the impact of customer and hotel types across different countries. Together, these visualizations provide a robust framework for understanding and optimizing decision-making in the hospitality sector.

For the hierarchical sunburst chart, the theoretical underpinnings from Luhmann's systems theory and Ostrom's institutional frameworks suggest that stable resource governance (in our case, market stability and pricing integrity) emerges when multi-level decision-making structures are visible and comprehensible. By enabling users to “drill down” from national policy contexts to hotel-level customer segments, this visualization concretely operationalizes the notion that institutional memory and adaptive capacity depend on transparent information flows across hierarchical levels. For example, if a country imposes a tourism tax, one can trace its ripple effects down to specific hotel types and customer categories, thereby illustrating how macro-level rules shape micro-level decisions in pricing and cancellation strategies.

Similarly, the geographic map extends the meta-organizational perspective by revealing how spatial clustering of ADR patterns and cancellation rates is not merely a statistical artifact but can be linked to regional governance forms and collective decision-making arrangements. For instance, areas with consistently high ADRs but low cancellation rates may reflect cooperative local tourism councils, strong regulatory frameworks, or stable partnership networks among stakeholders. By highlighting these spatial patterns, the map allows us to hypothesize that a region's policy climate—and the synergy between local

authorities, businesses, and community groups—may directly influence market outcomes. While additional ethnographic or policy–document analyses may be required to fully substantiate these hypotheses, the visualization itself provides a starting point for theoretically informed, empirically testable claims about the link between institutional structures and observed market patterns.

To strengthen the empirical grounding of our illustrative hospitality sector analysis, we incorporate quantitative indicators derived from industry datasets. For instance, aggregated booking records spanning 2015–2020 from major European hospitality platforms reveal that transient customers constituted approximately 82.6% of total reservations, while city hotels exhibited a 74.9% cancellation representation. Additionally, ADR fluctuations peaked during summer months, showing a median 15% increase from off-season baselines. These metrics were complemented by demographic and regional data to ensure comprehensive contextualization. By fusing these verified data sources with advanced visualization and modeling techniques, we enhance the reliability and generalizability of our policy-relevant insights.

5.7. Analysis Using Graph Neural Networks

To further enhance the framework’s predictive capabilities and directly address Luhmann’s [21] concept of system self-regulation, we implemented graph neural networks (GNNs) [34,35]. This implementation represents a concrete manifestation of how institutional structures can adapt and learn, aligning with Ostrom’s [19] principles of evolutionary institutional rules. Our GNN-based approach not only captures the complex relationships in the hospitality dataset but also demonstrates how modern AI systems can embody the institutional adaptation mechanisms discussed in our theoretical framework.

5.7.1. Network Architecture

The implemented architecture combines enhanced graph convolutional networks (GCNs) [36] with attention mechanisms:

- Multi-scale convolutions: three-layer GCN with an expanding–contracting structure.
- Feature transformation: parallel node feature processing through multi-layer perceptrons.
- Attention integration: edge attention mechanisms for relationship weighting.
- Residual connections: skip connections between consecutive layers.
- Batch normalization: layer-specific normalization for training stability.

5.7.2. Model Performance

The performance metrics of our GNN implementation provide an empirical validation of the theoretical principles outlined by Berkowitz and Grothe-Hammer [22] regarding adaptive decision-making structures. Cross-validation experiments across five folds revealed consistent performance that demonstrates the framework’s ability to maintain institutional trust through reliable outputs:

- Prediction accuracy: values of $73.4\% \pm 0.4\%$ across folds.
- Training convergence: stable loss reduction from 0.70 to 0.46.
- Model stability: small variance in performance metrics ($\sigma \approx 0.004$).
- Learning dynamics: effective early stopping around epoch 280–330.

To validate the robustness of our GNN implementation, we conducted a comprehensive five-fold cross-validation analysis, visualized in Figure 8. The training dynamics reveal several key insights into the model’s learning behavior. First, the validation loss shows consistent convergence patterns across all folds, starting from approximately 0.60 and stabilizing around 0.47, indicating robust generalization capabilities. Second, the

validation accuracy demonstrates stable improvement, reaching a plateau at approximately 73% across all folds with minimal variance ($\sigma \approx 0.004$). The training loss exhibits smooth decay characteristics, while the learning rate adaptation through the OneCycleLR scheduler shows effective exploration of the parameter space during early training followed by fine-tuning in later epochs. This example provides an instantiation of a dashboard for the control of advanced AI techniques.




Figure 8. Training dynamics across five-fold cross-validation of the GNN model. Interactive visualization: <https://api.wandb.ai/links/decurto-universidad-pontificia-comillas/kvhl87um>, accessed on 1 January 2025.

Figure 8 shows the training dynamics across five-fold cross-validation of the GNN model showing convergence patterns: (top row) validation loss, training epochs progression, and validation accuracy; (bottom row) training loss and learning rate adaptation. The consistent patterns across folds demonstrate the model’s stability, with validation accuracy converging around 73% and losses steadily decreasing across training steps. The learning rate curve shows effective adaptation through the OneCycleLR scheduler, supporting robust model optimization.

5.7.3. Architectural Insights

The enhanced GCN architecture revealed several key findings:

- Feature importance: room pricing (ADR) and lead time emerged as crucial predictors.
- Temporal patterns: strong seasonal effects captured through graph structure.
- Market segmentation: clear differentiation between city and resort hotel patterns.
- Regional effects: significant spatial correlations in cancellation behaviors.

The graph-based analysis provides additional validation of the theoretical framework’s ability to integrate several advanced techniques to capture complex relationships in hospitality data. The GNN implementation demonstrates the framework’s extensibility to advanced sophisticated approaches while maintaining interpretability through attention mechanisms and feature importance analysis.

5.7.4. Methodological Integration

Our choice of technical implementations—from visualization hierarchies to GNN architectures—was guided by the theoretical principles outlined in the framework. Each component was selected and designed to maintain the balance between institutional stability, as emphasized by Ostrom [19], and adaptive capacity, as discussed in contemporary organizational theory [20]. This deliberate alignment ensures that our technical solutions serve not just practical needs but also support the broader goals of institutional order and social cohesion.

5.8. Synthesis of Framework Components

The combination of visualization techniques, GNN implementation, and subsequent consideration of LLM integration demonstrates how our framework operationalizes the theoretical concepts of institutional adaptation and social order maintenance. This multi-layered approach embodies the process sequencing principles [6,7] while enabling the type of dynamic, context-aware decision-making envisioned in our theoretical framework. The progression from basic visualization through advanced AI techniques shows how modern technological capabilities can be harnessed while preserving institutional structures and social order mechanisms.

5.9. Integration of Large Language Models

The framework's capabilities can be significantly enhanced through the strategic integration of large language models (LLMs) [37,38] as a post-visualization analysis layer. This integration offers several key advantages:

5.9.1. Enhanced Pattern Analysis

LLMs can be leveraged to:

- Generate natural language explanations of complex visualization patterns.
- Identify non-obvious correlations across multiple visualization outputs.
- Provide context-aware interpretations of temporal and spatial trends.
- Synthesize insights from multiple visualization types into coherent narratives.

5.9.2. Interactive Decision Support

The framework can utilize LLMs to:

- Respond to natural language queries about visualization insights.
- Generate hypotheses about observed patterns and anomalies.
- Identify potential policy interventions based on identified trends.
- Provide real-time analysis of new data points within existing visualizations.

5.9.3. Knowledge Integration

LLMs can enhance the framework by:

- Connecting visualization insights with relevant academic literature and case studies.
- Identifying similar patterns from other sectors or geographic regions.
- Suggesting additional data sources or visualization approaches.
- Providing comparative analysis with historical trends or industry benchmarks.

5.9.4. Implementation Considerations

To effectively integrate LLMs into the visualization framework:

1. Establish clear protocols for LLM interaction with visualization outputs.
2. Develop standardized prompts for consistent analysis across distinct use cases.
3. Implement validation mechanisms for LLM-generated insights.
4. Create feedback loops between visualization updates and LLM analysis.
5. Ensure transparency in distinguishing between data-driven visualizations and LLM-generated interpretations.

This integration represents a significant advancement in the framework's analytical capabilities, enabling more sophisticated interpretation of visualization outputs while maintaining the rigor of data-driven analysis. The combination of structured visualization techniques with LLM-powered analysis provides a powerful tool for policy-makers and stakeholders to derive actionable insights from complex datasets.

5.9.5. Policy Generation Through LLMs

The implementation leverages OpenAI's GPT-4 model to generate structured policy recommendations as in [37]. Our code establishes a PolicyAdvisor class that processes metrics through carefully crafted prompts incorporating institutional theory frameworks. The coherence evaluation component ensures theoretical alignment through a secondary scoring mechanism, maintaining high standards for policy output quality.

This allows for following the policy sequence considering the temporal effect, as proposed by [6,8,9]. This temporal sequence views policies as past, present, and future, and the predictive framework takes this into account.

LLM Policy Analysis Prompt Template

System context: you are an expert policy advisor integrating institutional theory with data-driven insights.

Analysis request: analyze the following metrics for [REGION]: - Model performance: [METRICS]; - Feature importance: [FEATURES]; - Spatial patterns: [PATTERNS]; - Temporal trends: [TRENDS].

Theoretical framework: Consider: 1. Luhmann's institutional structures; 2. Ostrom's collective action principles; 3. Social order integration; 4. Actionable interventions.

Required output format: Return structured JSON with: - policy_recommendation; - theoretical_basis; - implementation_strategy; - expected_impact.

As exemplified, our framework leverages LLMs in a simplified use case to generate context-aware policy recommendations based on the analyzed metrics. The model demonstrated sophisticated understanding of institutional theory and market dynamics, as evidenced by this representative output:

LLM Policy Analysis Output (model: gpt-4)

Policy recommendation:

- Invest in quality enhancements and competitive pricing in high-demand regions.
- Develop targeted intervention strategies for rural areas.
- Leverage seasonality through dynamic pricing strategies.

Theoretical basis:

- Aligns with Luhmann's systems theory regarding environmental influences.
- Implements Ostrom's principles for resource management.
- Incorporates social integration mechanisms.

Implementation strategy:

- Regional business collaborations for peak season optimization.
- Rural area marketing campaigns and local partnerships.
- Seasonal pricing adjustments and incentive programs.

Expected impact:

- Increased peak season revenue in high-demand regions.
- Improved rural market performance.
- More balanced annual occupancy rates.

Model coherence score: 0.9

The model achieved a high coherence score, indicating strong alignment between theoretical foundations and practical recommendations. This implementation demonstrates

how LLMs can effectively bridge the gap between data analysis and actionable policy interventions while maintaining theoretical rigor.

5.10. Evaluation Framework for LLM Policy Recommendations

To enhance the clarity of our LLM-based analysis, we introduce a structured evaluation strategy for assessing the quality of LLM-generated recommendations. Table 2 outlines the evaluation criteria used to ensure theoretical consistency, empirical validity, and actionable insights.

Table 2. Evaluation framework for LLM policy recommendations.

Evaluation Criterion	Description
Theoretical Consistency	Aligns with institutional theory principles (Luhmann, Ostrom).
Empirical Validity	Correlates with historical trends and policy benchmarks.
Explainability	Generates human-readable rationales with cited sources.
Actionability	Translates insights into concrete policy recommendations.
Adaptability	Dynamically adjusts to new data inputs and stakeholder feedback.

As introduced in the previous section, using OpenAI’s programmatic API in Python for model GPT-4o, we utilized policy recommendation prompts with in-context structured institutional constraints. The model achieved a coherence score exceeding 0.9, ensuring alignment with institutional theory.

This framework could be further adapted into related policy domains as described in the following example.

Example: food security policy recommendation generation

- Input data: food security index, economic growth rate, agricultural yield data.
- Generated policy:

“To mitigate food security risks in Sub-Saharan Africa, prioritize decentralized food distribution hubs, enhancing local resilience. Modeled interventions suggest a 12% improvement in food availability over five years”.
- Evaluation: high coherence with Ostrom’s principles of decentralized governance.

This evaluation strategy ensures that LLM-driven policy recommendations remain both theoretically grounded and practically actionable, maintaining alignment with established institutional frameworks.

6. Additional Use Cases from Related Policy Domains

To demonstrate the broader applicability of our framework, we present two additional policy analysis implementations in complementary domains: global food security and educational development. These implementations showcase how our framework’s principles can be adapted to different policy contexts while maintaining the core emphasis on data-driven analysis and social order integration.

The implementation of our framework across food security and educational domains required careful data preparation aligned with institutional theory principles. Following Ostrom’s emphasis on multi-level analysis, we structured our data processing pipeline into institutional layers:

First, at the constitutional level, we integrated UNESCO Institute for Statistics (UIS) data for education metrics and the Global Food Security Index for security indicators.

At the collective choice level, we processed temporal patterns in both datasets, focusing on key transition periods (2019–2022 for food security; 2000–2020 for education).

This temporal structuring aligns with Luhmann’s concept of institutional memory, allowing us to trace policy evolution through quantifiable metrics. The food security data preparation revealed three distinct institutional dimensions—affordability, availability, and quality/safety—which were preprocessed using pandas and integrated into interactive visualizations.

At the operational level, we developed standardized data pipelines using Python, focusing on institutional interoperability. This included creating long-format datasets suitable for temporal analysis and geospatial visualization, with careful attention being paid to missing value imputation and validation of metrics across institutional sources. The resulting preprocessed datasets enabled the development of interactive visualizations that maintain institutional integrity while supporting dynamic policy analysis.

This methodological approach demonstrates how institutional theory can guide not only analysis but also the fundamental organization and preparation of policy-relevant data. The following subsections detail the specific implementations and insights derived from this theoretically grounded data preparation process.

6.1. Exploratory Data Analysis

The development of our interactive tools was informed by an initial exploratory analysis using Tableau visualizations (Figure 9). These preliminary visualizations helped identify key patterns and relationships in the data, guiding the development of more sophisticated interactive tools. The food security choropleth map revealed global patterns in security metrics, while the literacy rates visualization highlighted specific challenges in sub-Saharan Africa. These insights were then incorporated into the development of more comprehensive analytical tools using Streamlit (Figures 10 and 11).



Figure 9. Left: Global food security choropleth map showing worldwide distribution of security metrics. Interactive version available at: <https://public.tableau.com/app/profile/decurto/viz/GlobalFoodSecurityComparingKeyIndicatorsAcrossCountries/Sheet1>, accessed on 1 January 2025. Right: Adult literacy rates visualization, focusing on sub-Saharan Africa. Interactive version available at: <https://public.tableau.com/app/profile/dezarza/viz/AdultLiteracyRatesinSub-SaharanAfrica/Sheet1>, accessed on 1 January 2025.

6.2. Global Food Security Analysis

The first implementation focuses on global food security metrics across countries, providing comparative analysis through interactive visualizations. The implementation combines choropleth mapping for spatial analysis with detailed metric comparisons using Streamlit; see Figure 10. This visualization story enables policy-makers to:

- Compare food security scores across different geographical regions.
- Track temporal changes in food security metrics between 2019 and 2022.
- Analyze the relationship between different food security dimensions through radar charts.
- Identify leading countries and potential policy models through ranked comparisons.




Figure 10. Visualization dashboard of global food security showing comparative metrics across countries. Interactive version available at: <https://foodsecurity-decurto.streamlit.app/>, accessed on 1 January 2025.




Figure 11. Dashboard of educational development tracking, focusing on literacy rates. Interactive version available at: <https://gloaleducation-dezarza.streamlit.app/>, accessed on 1 January 2025.

In particular, our analysis of the Global Food Security Index dataset through interactive visualizations revealed several significant patterns and policy-relevant insights. The comparative analysis of food security scores between 2019 and 2022 highlighted notable regional variations and the impact of global events. For instance, countries like Finland and Spain demonstrated marked improvements in their food security metrics, with Spain showing progress in affordability metrics, suggesting successful economic policies enhancing purchasing power. In contrast, countries like Germany and Singapore experienced significant declines, albeit for different reasons. Germany's decline was primarily driven by reduced availability scores, potentially reflecting supply chain disruptions from the Ukraine conflict, while Singapore's decrease was concentrated in quality and safety metrics, possibly due to shifts in food import sources amid global trade tensions.

The ranking analysis of top-performing countries reveals a strong Nordic presence, with Finland, Ireland, and Norway leading the global rankings. This clustering suggests the effectiveness of regional policy approaches and the potential benefits of policy transfer within geographic regions. The radar chart analysis provides deeper insights into these performance patterns. For example, while some countries maintained relatively stable overall scores, the underlying dimensional changes tell a more complex story. The visualization reveals how trade policies, geopolitical events, and economic conditions differently affect the three key dimensions of food security—affordability, availability, and quality/safety.

These findings demonstrate the framework’s capability to capture and visualize complex policy interactions. The COVID-19 pandemic and Ukraine conflict emerge as significant external shocks with varying regional impacts, highlighting the importance of resilient food security systems. Countries showing improved metrics despite these challenges offer valuable policy lessons.

6.3. Educational Development Tracking

The second implementation focuses on educational metrics, particularly adult literacy rates in sub-Saharan Africa, demonstrating the framework’s utility in tracking development goals. Through interactive visualizations (see Figure 11), this implementation enables:

- Monitoring progress toward UNESCO’s adult literacy targets.
- Comparing educational outcomes across different regions.
- Analyzing temporal trends in literacy rates.
- Identifying successful policy interventions through comparative analysis.

Our analysis of the UNESCO Institute for Statistics (UIS) dataset through interactive visualizations revealed striking patterns in educational development across regions. The comparative analysis between sub-Saharan Africa and South America proved particularly insightful, highlighting both success stories and persistent challenges. In Latin America, the data demonstrates remarkable progress in primary education completion rates, with countries like Chile, Uruguay, and Ecuador approaching 100% completion rates between 2000 and 2020. Similarly, positive trends are observed in adult literacy rates, with Uruguay maintaining consistently high performance, and secondary education showing steady improvements across the region, though with continuing room for progress.

The sub-Saharan African data reveal more complex patterns, characterized by significant inter-country disparities across all three key indicators. South Africa emerges as a regional leader, while Kenya shows substantial progressive improvement over time. However, many countries in the region continue to face challenges in reaching international benchmarks. This variance in outcomes suggests the importance of understanding local contexts in policy implementation. The introduction of a 90% target line for adult literacy rates provides particularly striking visualization of these regional disparities: while some countries show steady progress toward this goal, others like Ethiopia face persistent challenges in approaching this benchmark.

The interactive capabilities of our framework proved especially valuable in exploring these educational metrics. The ability to dynamically compare regions, countries, and time periods revealed both broad trends and specific policy outcomes. For instance, the visualization clearly demonstrates how consistent educational policies in Latin America have yielded sustained improvements, while the more varied approaches in sub-Saharan Africa have resulted in more diverse outcomes. These insights suggest the potential value of policy transfer between regions while acknowledging the need for local adaptation.

The framework’s capacity to handle multilingual data sources and complex temporal trends proved particularly valuable in the education sector, where data collection and reporting practices vary significantly across regions. The integration of multiple

indicators—from basic literacy rates to detailed enrollment statistics—provides a comprehensive view of educational development that can inform policy decisions at multiple levels, from local interventions to international development strategies.

These additional implementations demonstrate how our framework's principles of data-driven analysis, interactive visualization, and social order integration can be effectively applied across different policy domains. Both cases utilize similar technical approaches—combining Tableau for initial data exploration with Streamlit for interactive deployment—while adapting the specific metrics and visualizations to their respective domains. This adaptability supports our framework's potential for broader application in policy analysis and intervention design.

These visualizations could be further enhanced through real-time integration with LLMs like Google's Gemini 2.0, with text and visual capabilities (<https://aistudio.google.com/live>, accessed on 1 January 2025; model: gemini2.0-flash-exp) that allow for direct screen sharing with the AI agent. Such an integration would enable interactive man-machine analysis of the visualizations, where the model could:

- Provide real-time natural language explanations of observed patterns.
- Suggest potential policy interventions based on identified trends.
- Compare current metrics with historical data to identify significant changes.
- Offer cross-domain insights by analyzing relationships between food security and educational development metrics.
- Generate automated reports and summaries for policy-makers.

For example, when analyzing the Global Food Security visualization, Gemini could identify correlations between different regions' security metrics and suggest targeted interventions based on successful policies in similar contexts. Similarly, for the adult literacy rates visualization, the model could provide insights into effective educational strategies by analyzing temporal patterns and regional variations. This human-AI collaborative approach would enhance the framework's analytical capabilities while maintaining the interpretability and actionability of insights for policy-makers.

The integration of real-time AI analysis with interactive visualizations exemplifies our framework's core theoretical principles of institutional adaptation and social order maintenance. Following Luhmann's [21] concept of systems maintaining integrity through communication, these AI-enhanced visualizations serve as institutional structures that facilitate structured dialogue between data, algorithms, and human decision-makers. The combination of Tableau dashboards, Streamlit applications, and Gemini's analytical capabilities creates what Ahrne and Brunsson [20] would recognize as a partial organization—a decentralized yet coherent system for knowledge generation and policy formation. This technological ensemble embodies Ostrom's [19] principles of institutional adaptation, where the interaction between human expertise and AI analysis creates dynamic, self-regulating mechanisms for policy development. The framework thus demonstrates how modern computational tools can be leveraged to strengthen, rather than disrupt, institutional structures while maintaining social order through transparent, interpretable, and theoretically grounded analysis.

7. Discussion

The proposed framework demonstrates an innovative approach to integrating data-driven methodologies for policy analysis and intervention design, particularly exemplified through its application to the hospitality sector. We present a detailed analysis connecting the key elements of our research.

The initial framework (as depicted in Figures 1 and 2) establishes five critical components: data collection, historical analysis, policy impact assessment, predictive modeling,

and intervention design. This theoretical structure is demonstrated through the visualizations presented in Figures 5–7, as well as in Figures 9–11, which reveal how data can be analyzed and presented across multiple dimensions.

For instance, the ADR interactive map visualization (see Figure 5) directly supports the framework's data collection and historical analysis components by providing geographical and temporal insights into pricing patterns. The violin plot comparison between city and resort hotels enhances policy impact assessment through its revelation of distinct pricing behaviors and market segmentation. Additionally, the temporal patterns visualization strengthens predictive modeling capabilities by illuminating clear seasonal trends and patterns.

The communication structure (as illustrated in Figures 3 and 4) emphasizes three layers—determination, efficiency, and interaction—each reflected distinctly in the visualization implementation. The determination layer manifests through a structured presentation of data across different visualization types, while the efficiency layer is particularly evident in the information density achieved, for example, through the sunburst visualization (Figure 7). The interaction layer comes to life through the interactive features woven throughout all three visualization stories.

The implementation of our framework across different policy domains—from hospitality to food security and educational development—demonstrates its versatility and theoretical robustness. Each domain showcases how the framework's five core components adapt to different contexts while maintaining institutional integrity. In the food security analysis, the data collection and historical analysis components revealed how global events like COVID-19 and geopolitical conflicts impact institutional structures differently across regions. Similarly, in educational development tracking, the policy impact assessment and predictive modeling components helped identify successful institutional adaptations, particularly evident in Latin America's consistent educational improvements. These additional implementations demonstrate how Luhmann's [21] concept of system integrity through communication manifests across different policy domains—from hotel booking patterns to food security metrics and educational outcomes. The framework's efficiency layer, particularly evident in the interactive visualizations, enables what Ahrne and Brunsson [20] describe as partial organization, allowing for decentralized yet coherent analysis across diverse institutional contexts. This is particularly evident in how the framework handles multi-dimensional data: from ADR patterns in hospitality to regional food security metrics and educational development indicators. Moreover, following Ostrom's [19] principles of institutional adaptation, these implementations show how the framework can maintain social order while facilitating policy learning across domains—for instance, how successful institutional structures in educational development might inform approaches to food security governance, or how patterns of institutional adaptation in the hospitality sector might offer insights for educational policy design.

While different tools were used across the hospitality, food security, and education domains, this selection was intentional to optimize the analytical insights for each domain. The hospitality sector analysis leveraged Tableau and Flourish for interactive data storytelling, given the structured nature of booking data. For food security, Streamlit dashboards enabled policy-makers to engage in comparative metric evaluation. The selection of tools was driven by domain-specific data requirements, visualization complexity, and the need for AI-based predictive analytics, ensuring an optimal balance of interpretability and computational depth.

This use of advanced visualization techniques, LLMs, and institutional theory demonstrates the framework's capacity for dynamic adaptation while preserving social order mechanisms. By combining Streamlit's interactive capabilities with, for instance, Gemini's

analytical power and grounding in Ostrom's institutional principles, we create a robust system for evidence-based policy-making. The framework's success across food security metrics and educational development indicators suggests broader applicability to other complex policy domains. Moreover, the complementary strengths of visualization tools and LLMs create what we term "institutional intelligence"—a hybrid approach where human expertise, AI capabilities, and theoretical foundations work in concert to enhance policy analysis while maintaining institutional integrity. This synthesis of technical innovation and theoretical rigor offers a blueprint for future policy analysis systems that can adapt to emerging challenges while preserving essential institutional structures and social order mechanisms.

8. Conclusions

This study presents a comprehensive framework for integrating data-driven approaches into policy analysis and intervention design, demonstrated through its application to the hospitality sector. Our approach offers two key theoretical contributions: First, by embedding institutional theories of social order and collective action within a data-driven analytical framework, we demonstrate that concepts like organizational adaptability, institutional memory, and communication structures can be operationalized quantitatively. Second, we illustrate how these theoretical constructs provide interpretive depth to empirical findings, showing how interdependent variables are outcomes influenced by layers of governance, stakeholder negotiations, and historical institutional arrangements.

Moreover, our research introduces the concept of "institutional intelligence"—the synergistic integration of human expertise, AI capabilities, and institutional theory. This framework demonstrates how institutional intelligence can emerge through the combination of LLMs, interactive visualizations, and domain knowledge, creating a hybrid system that preserves institutional memory while enabling dynamic adaptation. The resulting system not only processes data but also understands and maintains institutional contexts, offering a new paradigm for policy analysis that balances innovation with institutional stability. Future work could explore how institutional intelligence might evolve through advanced LLM architectures, enhanced visualization capabilities, and deeper integration with domain-specific knowledge bases, potentially leading to more sophisticated forms of human–AI collaboration in policy-making. This advancement in institutional intelligence represents a significant step toward more sophisticated, context-aware policy analysis systems that can effectively bridge the gap between theoretical frameworks and practical implementation challenges.

This synthesis reveals that complex policy environments cannot be adequately understood through data or theory alone. Instead, the productive tension between theoretical models of institutional resilience and empirical evidence generates actionable insights for sustainable, inclusive policy-making. The framework's success in combining multiple data analysis approaches, enabling effective communication through interactive visualizations, and establishing robust theoretical foundations demonstrates its potential for advancing evidence-based policy formation across diverse sectors.

Several promising directions emerge for future research. First, while our framework has proven effective in the hospitality sector, in global education, and in food security, its application to other domains such as healthcare policy, environmental regulation, or urban planning could yield valuable insights about its generalizability and domain-specific adaptations. Second, the integration of more sophisticated LLM architectures, the proper use of retrieval-augmented generation (RAG), and fine-tuning approaches could enhance the framework's ability to generate contextually appropriate policy recommendations while maintaining theoretical rigor.

Further research could also explore the development of more advanced visualization techniques that can handle higher-dimensional policy spaces and complex stakeholder relationships. Additionally, investigating the framework's performance with real-time data streams and dynamic policy environments would enhance its practical utility. The incorporation of causal inference methods could strengthen the framework's ability to identify policy intervention effects and their propagation through social systems.

Another important direction involves expanding the framework's capability to handle multi-stakeholder policy optimization, potentially through the integration of game-theoretic approaches with deep learning architectures. Finally, developing standardized metrics for evaluating the long-term effectiveness of policy interventions generated through our framework would provide valuable feedback for continuous improvement and adaptation.

Author Contributions: Conceptualization, J.d.C., I.d.Z. and V.S.-F.; data curation, J.d.C. and I.d.Z.; formal analysis, J.d.C., I.d.Z. and L.S.F.; funding acquisition, C.T.C. and V.S.-F.; investigation, J.d.C., I.d.Z. and L.S.F.; methodology, J.d.C. and I.d.Z.; software, J.d.C. and I.d.Z.; supervision, V.S.-F. and C.T.C.; validation, L.S.F. and V.S.-F.; visualization, J.d.C. and I.d.Z.; writing—original draft, J.d.C. and I.d.Z.; writing—review and editing, all authors. All authors have read and agreed to the published version of the manuscript.

Funding: The research received the funding from the BARCELONA Supercomputing Center (TIFON). Leandro Fervier is supported by “Convocatoria de Ayudas para Iberoamericanos en Estudios de Doctorado from Universidad de Zaragoza and Santander Universidades”.

Data Availability Statement: All data underlying the results presented in this study are publicly available. The datasets used in this study can be accessed from the following repositories: the hospitality sector data are available from the Data in Brief repository (<https://doi.org/10.1016/j.dib.2018.11.126>); the Global Food Security Index data are available from The Economist Impact (<https://impact.economist.com/sustainability/project/food-security-index>) and from Kaggle (<https://www.kaggle.com/datasets/rahuljangam/global-food-security-index>); and educational statistics are available from the UNESCO Institute for Statistics (<https://uis.unesco.org/bdds>). Additional processed data and interactive visualizations supporting our findings are accessible through the following public links: the ADR visualization story and associated data are available on the Flourish platform (<https://public.flourish.studio/story/2733675/>, accessed on 1 January 2025); the complete visualization sequence with booking patterns and cancellation rates is available on Tableau Public (<https://public.tableau.com/app/profile/decurnto/viz/Tendnciesdereservesdhotelsilescancellacions/Story1>, accessed on 1 January 2025); the interactive sunburst visualization implementation and source code are hosted at <https://decurnto01.netlify.app/> (accessed on 1 January 2025); the GNN training metrics, model performance data, and visualization outputs are available via the Weights & Biases platform (<https://api.wandb.ai/links/decurnto-universidad-pontificia-comillas/kvhl87um>, accessed on 1 January 2025); the Food Security dashboard is accessible at <https://foodsecurity-decurnto.streamlit.app/> (accessed on 1 January 2025); and the Educational Development dashboard is available at <https://globaleducation-dezarza.streamlit.app/> (accessed on 1 January 2025). Visualizations for Food Security (<https://public.tableau.com/app/profile/decurnto/viz/GlobalFoodSecurityComparingKeyIndicatorsAcrossCountries/Sheet1>, accessed on 1 January 2025) and Educational Development (<https://public.tableau.com/app/profile/dezarza/viz/AdultLiteracyRatesinSub-SaharanAfrica/Sheet1>, accessed on 1 January 2025) are also hosted on Tableau Public. All other relevant data are contained within the manuscript, and the visualizations along with their underlying data will be preserved at the specified links with unrestricted public access.

Acknowledgments: During the preparation of this work, the authors used ChatGPT/Copilot/Gemini/Claude for code completion and pseudocode generation, Grammarly for English language editing, and Mendeley for bibliography management. The authors reviewed and edited the content as needed and take full responsibility for the content of this publication. The authors would like to thank the

BARCELONA Supercomputing Center for providing access to MareNostrum 5 and technical support throughout this research.

Conflicts of Interest: The authors declare that they have no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial intelligence
LLMs	Large language models
GNNs	Graph neural networks
GCNs	Graph convolutional networks
ADR	Average daily rate

References

1. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models are Few-Shot Learners. In Proceedings of the Advances in Neural Information Processing Systems, Online, 6–12 December 2020; Volume 33, pp. 1877–1901.
2. De Zarzà, I.; de Curtò, J.; Roig, G.; Manzoni, P.; Calafate, C.T. Emergent Cooperation and Strategy Adaptation in Multi-Agent Systems: An Extended Coevolutionary Theory with LLMs. *Electronics* **2023**, *12*, 2722. [\[CrossRef\]](#)
3. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; Volume 30, pp. 5998–6008.
4. Grothe-Hammer, M.; Berkowitz, H. Unpacking Social Order: Toward a Novel Framework That Goes Beyond Organizations, Institutions, and Networks. *Crit. Sociol.* **2024**, *50*, 1399–1420. [\[CrossRef\]](#)
5. Pierson, P. When Effect Becomes Cause: Policy Feedback and Political Change. *World Politics* **1993**, *45*, 595–628. [\[CrossRef\]](#)
6. Daugbjerg, C. Sequencing in public policy: The evolution of the CAP over a decade. *J. Eur. Public Policy* **2009**, *16*, 395–411. [\[CrossRef\]](#)
7. Howlett, M. Process sequencing policy dynamics: Beyond homeostasis and path dependency. *J. Public Policy* **2009**, *29*, 241–262. [\[CrossRef\]](#)
8. Kay, A.; Baker, P. What can causal process tracing offer to policy studies? A review of the literature. *Policy Stud. J.* **2015**, *43*, 1–21. [\[CrossRef\]](#)
9. Linder, S.H.; Peters, B.G. A design perspective on policy implementation: The fallacies of misplaced prescription. *Rev. Policy Res.* **1987**, *6*, 459–475. [\[CrossRef\]](#)
10. Milán Paradela, T. ¿El marco de diseño de políticas o diseñando un marco para el análisis de políticas? *Cuad. Cendes* **2019**, *36*, 87–115.
11. Rodrigues de Caires, C. Peters, B.G. Policy Problems and Policy Design. Cheltenham: Edward Elgar. *Cienc. Política* **2019**, *14*, 315–321. [\[CrossRef\]](#)
12. Alison, G.T. Modelos conceptuales y la crisis de los misiles cubanos. In *La Hechura de las Políticas*; Grupo Editorial Porrúa: Mexico City, Mexico, 1992; pp. 119–200.
13. Lindblom, C. *El Proceso de Elaboración de Políticas Públicas*, 1st ed.; Grupo Editorial Miguel Ángel Porrúa: Mexico City, Mexico, 1991.
14. Bachrach, P.; Baratz, M.S. Decisions and not decisions: An analytical framework. *Gestión Estrateg.* **2009**, *35*, 81–95. [\[CrossRef\]](#)
15. Halbwachs, M. *On Collective Memory*; University of Chicago Press: Chicago, IL, USA, 2020.
16. Schneider, A.; Ingram, H. Systematically pinching ideas: A comparative approach to policy design. *J. Public Policy* **1988**, *8*, 61–80. [\[CrossRef\]](#)
17. Teece, D.J. Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strateg. Manag. J.* **2007**, *28*, 1319–1350. [\[CrossRef\]](#)
18. Eisenhardt, K.M.; Martin, J.A. Dynamic capabilities: What are they? *Strateg. Manag. J.* **2000**, *21*, 1105–1121. [\[CrossRef\]](#)
19. Ostrom, E. *Governing the Commons: The Evolution of Institutions for Collective Action*; Cambridge University Press: Cambridge, MA, USA, 1990.
20. Ahrne, G.; Brunsson, N. Organization Outside Organizations: The Significance of Partial Organization. *Organization* **2011**, *18*, 83–104. [\[CrossRef\]](#)

21. Luhmann, N. *Sistemas Sociales: Lineamientos para una Teoría General*; Antrophos Editorial: Sankt Augustin, Germany, 1998.
22. Berkowitz, H.; Grothe-Hammer, M. From a Clash of Social Orders to a Loss of Decidability in Meta-Organizations Tackling Grand Challenges: The Case of Japan Leaving the International Whaling Commission. *Res. Sociol. Organ.* **2022**, *79*, 115–138. [[CrossRef](#)]
23. Scartascini, C.; Spiller, P.; Stein, E.; Tommasi, M. *El Juego Político en América Latina: ¿Cómo se Deciden las Políticas Públicas?* Banco Interamericano de Desarrollo: Washington, DC, USA, 2010.
24. Walker, W.E. Policy analysis: A systematic approach to supporting policymaking in the public sector. *J. -Multi-Criteria Decis. Anal.* **2000**, *9*, 11–27. [[CrossRef](#)]
25. Papadopoulos, P.; Soflano, M.; Chaudy, Y.; Adejo, W.; Connolly, T.M. A systematic review of technologies and standards used in the development of rule-based clinical decision support systems. *Health Technol.* **2022**, *12*, 713–727. [[CrossRef](#)]
26. Sampson, C.; Zamora, B.; Watson, S.; Cairns, J.; Chalkidou, K.; Cubi-Molla, P.; Devlin, N.; García-Lorenzo, B.; Hughes, D.A.; Leech, A.A.; et al. Supply-side cost-effectiveness thresholds: Questions for evidence-based policy. *Appl. Health Econ. Health Policy* **2022**, *20*, 651–667. [[CrossRef](#)]
27. Luján, J.L. Evidence-based policies: Lessons from regulatory science. *Politics Policy* **2023**, *51*, 524–537. [[CrossRef](#)]
28. Newman, J.; Mintrom, M. Mapping the discourse on evidence-based policy, artificial intelligence, and the ethical practice of policy analysis. *J. Eur. Public Policy* **2023**, *30*, 1839–1859. [[CrossRef](#)]
29. Narne, S.; Adedjoja, T.; Mohan, M.; Ayyalasomayajula, T. AI-Driven Decision Support Systems in Management: Enhancing Strategic Planning and Execution. *Int. J. Recent Innov. Trends Comput. Commun.* **2024**, *12*, 268–276.
30. Safaei, M.; Longo, J. The End of the Policy Analyst? Testing the Capability of Artificial Intelligence to Generate Plausible, Persuasive, and Useful Policy Analysis. *Digit. Gov. Res. Pract.* **2024**, *5*, 1–35. [[CrossRef](#)]
31. Liao, Z.; Wang, J.; Shi, Z.; Lu, L.; Tabata, H. Revolutionary potential of ChatGPT in constructing intelligent clinical decision support systems. *Ann. Biomed. Eng.* **2024**, *52*, 125–129. [[CrossRef](#)]
32. Sanagustín Fons, M.V.; Moseñe Fierro, J.A.; Gómez y Patiño, M. Rural tourism: A sustainable alternative. *Appl. Energy* **2011**, *88*, 551–557. [[CrossRef](#)]
33. Antonio, N.; de Almeida, A.; Nunes, L. Hotel booking demand datasets. *Data Brief* **2019**, *22*, 41–49. [[CrossRef](#)]
34. Kipf, T.N.; Welling, M. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv* **2017**, arXiv:1609.02907.
35. Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; Bengio, Y. Graph Attention Networks. *arXiv* **2018**, arXiv:1710.10903.
36. Zhang, S.; Tong, H.; Xu, J.; Maciejewski, R. Graph convolutional networks: A comprehensive review. *Comput. Soc. Netw.* **2019**, *6*, 1–23. [[CrossRef](#)]
37. De Zarzà, I.; de Curtò, J.; Roig, G.; Calafate, C.T. Optimized Financial Planning: Integrating Individual and Cooperative Budgeting Models with LLM Recommendations. *AI* **2024**, *5*, 91–114. [[CrossRef](#)]
38. Bansal, G.; Wortman Vaughan, J.; Amershi, S.; Horvitz, E.; Fourney, A.; Mozannar, H.; Dibia, V.; Weld, D.S. *Challenges in Human-Agent Communication*; Technical Report MSR-TR-2024-53; Microsoft: Washington, DC, USA, 2024.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.