Proceedings of the 3rd International Conference on Machine Learning and Automation
DOI: 10.54254/2755-2721/2025.29779

Market Response Estimation via Causal Representation Learning on Film Policy Texts

Xinye Xu[1*], Xuanqing Huang[2]

[1]*University of Westminster, London, United Kingdom*
[2]*School of the Art Institute of Chicago, Chicago, USA*
**Corresponding Author. Email: rara481846778@gmail.com*

Abstract. This paper introduces a causal representation learning model to quantitatively predict the impact of film policy texts on the observable market response of China's film industry. By integrating semantic embedding, structural causal models, and counterfactual simulation, this model establishes a stable correspondence between policy rhetoric and behavioral patterns. This framework has revealed the hidden causal connection between language structure and production scale, investment, and box office. Experimental analysis shows that the prediction accuracy is quite considerable, with the mean absolute error (MAE) being 0.183±0.012 and the causal effect stability (CES) reaching 0.924±0.006, which is over 20% higher than the baseline value. This framework also provides a causal explanation. It shows that the policy wording oriented towards subsidies and innovation has a significant positive impact ($\beta = +0.47\pm0.05, p < 0.01$), while the restrictive content control clause has an adverse impact ($\beta = -0.36 \pm 0.04, p < 0.05$). The prediction results obtained from the counterfactual analysis indicate that Adjusting restrictive clauses to innovation-oriented ones is expected to drive investment growth by 14.8±2.1%. Policy rhetoric may be a signal that can measure the economy, connecting computational linguistics and film market analysis with evidence-based cultural policies.

Keywords: Causal Representation Learning, Policy Text Analysis, Film Industry, Market Response Estimation

1. Introduction

The mechanism by which policy text content has a quantifiable impact on the economy remains a blind spot to be explored in the research field. Traditional econometric models treat policies as binary events or dummy variables, ignoring extensive semantic and causal heterogeneity, which is an attribute of policy talks [1]. Within the realm of cultural economy, film policy is an extremely crucial governance tool. It plays a decisive role in the flow of investment, production models, and the viewing conditions of audiences. Policy influences the distribution pattern of monetary resources, defines the scope of market access, and also sets an upper limit for creativity [2].

In practical application, the language, style and composition priority of policy documents will implicitly reflect investor confidence and industry coordination [3]. Policies focusing on "subsidy enhancement" will stimulate active financing and promote project initiation, while policies

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2025.29779

emphasizing "content review" will induce hesitation about risks and delay the production cycle. To understand this calculation method, It is necessary to link language semantics with causal reasoning.

This study constructs a causal Representation Learning (CRL) framework, which can capture the semantics of policy texts that trigger measurable market responses [4].

This framework combines domain-tuned language models with structural causal diagrams and counterfactual simulations. Unlike sentiment analysis and topic modeling, this idea predicts causal semantics rather than correlation features. It helps researchers and policymakers simulate hypothetical policy shifts and thereby estimate their related economic impacts.

The main contributions of this work are reflected in three dimensions. It enables the causal mapping from "language to economy" to be coordinated, embedding policy semantics into an interpretable causal space. The integration of computational linguistics, causal reasoning and cultural economics provides a data-intensive paradigm for evaluating the effectiveness of governance in the creative sector.

2. Literature review

2.1. Policy text analysis in cultural economics

The advent of natural language processing (NLP) brought about automated semantic modeling [5], facilitating large-scale extraction of strategic intentions. In the early cultural field, early policy text analysis was qualitative evaluated through manual coding, which restricted scalability and objectivity. The vast majority of NLP applications only have correlations but lack causal explainability. The latest progress in embedded policy analysis shows that policy behavior is influenced.

2.2. Market response modeling

The language composition of policies has the opportunity to change producers' expected ideas. The traditional market response models in the film industry adopt regression elasticity and diffusion models. Models designed by scholars mainly rely on numerical indicators such as ticket prices, the number of screens, and audience attributes. In these models, there is generally no textual policy environment that can convey investor sentiment and resource allocation information [6]. Moreover, the distinction of text-based causal models has an impact on investment decisions long before measurable market activities emerge, which limits the explanatory power of existing methods.

2.3. Causal representation learning

Causal representation learning connects causal reasoning with deep learning by separating the unchanging causes that lead to observable results. It supports the estimation of counterfactual results and therapeutic effects through the cross-domain integration of representation learning and structural causal models. When this principle is extended to the scope of policy texts, it can be perceived how the causal structure of language affects economic signals in terms of language. This transforms the descriptive analysis of the text into an explanatory policy analysis [7].

DOI: 10.54254/2755-2721/2025.29779

3. Methodology

3.1. Causal representation framework

The model has a three-level composition:

A transformer-tuned domain represents each policy document as a dense semantic vector, while retaining contextual dependencies and regulatory purpose.

Structural connections between latent representations, mediating factors, and market responses are represented using a directed acyclic graph (DAG), written as follows [8]:

Equation (1)

$$Y = f(\mathbf{z}, \mathbf{m}, \epsilon), \quad \text{where } \mathbf{m} = g(\mathbf{z}, \eta) \tag{1}$$

Here, Y denotes market outcomes, $\mathbf{m}$ mediators (e.g., investor sentiment), and ϵ represent exogenous shocks.

Counterfactual inference predicts hypothetical outcomes under alternative policy wording via structural interventions as shown in Equation (2) [9]:

$$\mathrm{CF}(Y) = E\left[Y \mid \mathrm{do}\left(\mathbf{z} = \mathbf{z}^{l}\right)\right] \tag{2}$$

This modular structure supports both predictive learning and interpretive analysis, enabling dynamic simulation of market reactions.

3.2. Data construction and feature alignment

A dataset of 3,000 policy documents obtained from national and provincial regulations, covering a time span from 2010 to 2025, has been collated. Each document has marked the industries (production, distribution, financing) and types of penalties (subsidies, restrictions, innovation). The measurement indicators for the auxiliary market include box office revenue (in billions of yuan), the number of films, and investment situations. The consistency of the quarterly time ensures that the policy emerges before the observable results, while controlling the simultanoity bias. The final dataset includes the duration of each auxiliary variable for 60 quarters, and a 3000×60×8 tensor is available for modeling.

3.3. Training and optimization

The model jointly optimizes predictive and causal stability objectives as shown in Equation (3):

$$\mathscr{L} = \mathscr{L}_{pred} + \lambda \mathscr{L}_{inv} \tag{3}$$

where $\mathscr{L}_{pred}$ is mean squared error loss between predicted and actual outcomes, and $\mathscr{L}_{inv}$ enforces invariance of causal effects across temporal splits [10]. It employs a learning rate of 2e−5, with a batch size of 32, and combines cosine decay with the AdamW optimizer. It also covers early termination when the causal loss of the validation set does not improve for 10 consecutive epochs. It is achieved through PyTorch and DoWhy and runs on NVIDIA A100 Gpus.

DOI: 10.54254/2755-2721/2025.29779

4. Experimental design

4.1. Dataset partitioning and evaluation metrics

Evaluation metrics include MAE, RMSE, Average Treatment Effect (ATE), and Causal Effect Stability (CES). Statistical reliability is measured via bootstrap resampling (n = 1000).

Data are partitioned chronologically: training (2010–2019), validation (2020–2022), and testing (2023–2025).

4.2. Baseline comparisons and ablation studies

OLS Regression: linear mapping from policy features to outcomes. LSTM: sequential modeling of temporal dependencies. BERT Regression: contextual embeddings with regression output.

Causal Forest: non-parametric causal effect estimation. Ablation studies remove the counterfactual simulator and invariant loss terms to assess their contribution (see table 1).

Table 1. Summarizes the quantitative performance

Model	MAE (↓)	RMSE (↓)	ATE Stability (↑)	CES (↑)
OLS Regression	0.264 ± 0.018	0.392 ± 0.025	0.812 ± 0.017	0.798 ± 0.015
LSTM	0.231 ± 0.014	0.357 ± 0.022	0.845 ± 0.013	0.861 ± 0.009
BERT Regression	0.218 ± 0.013	0.349 ± 0.019	0.862 ± 0.010	0.877 ± 0.008
Causal Forest	0.206 ± 0.012	0.331 ± 0.018	0.879 ± 0.009	0.898 ± 0.007
CRL (ours)	0.183 ± 0.012	0.301 ± 0.016	0.905 ± 0.007	0.924 ± 0.006

The proposed CRL framework yields significant improvement ($p < 0.01$, paired t-test), particularly in ATE stability and causal robustness under domain shifts.

4.3. Implementation details and computational setup

After 50 experiments, the typical convergence time was 38±3 epoch, the average GPU memory usage reached 17.4 GB, and the training duration for each run was approximately 11 hours. The verification monitoring requirements were that the ATE variance was less than 0.005 and the loss stability slope was less than 0.002. Monte Carlo was used to conduct cross-verification on random splitting. It was confirmed that the generalization consistency of the model was ±2.7%.

5. Results and discussion

5.1. Quantitative results and performance evaluation

The CRL model achieved a statistically significant 17% reduction in MAE. The RMSE value decreased from 0.331±0.018 to 0.301±0.016. When there was a time shift, the average prediction error could be limited to ±3.1% or lower, demonstrating robustness. The model also showed a stable causal relationship: In the regional variation test, CES increased by 2.6%, while the variance of ATE decreased from 0.041 to 0.027. The average processing β of subsidy-oriented language was +0.47 plus or minus 0.05, and the average processing β of restrictive clauses was -0.36, with an error of ±0.04.

DOI: 10.54254/2755-2721/2025.29779

The average elasticity of the 12 policy areas is, which indicates that when external shocks remain unchanged, every 1% increase in positive policy expression can drive a 1.27% increase in market activity.

5.2. Interpretability and causal pathway analysis

Linguistic causal decomposition identified three dominant positive mediators:

(1) Fiscal subsidy intensity, contributing +0.41 ± 0.06 effect units;

(2) Co-production encouragement, +0.32 ± 0.05;

(3) Infrastructure support, +0.29 ± 0.04.

Conversely, content restriction (−0.37 ± 0.05) and import control (−0.34 ± 0.06) acted as negative pathways (see table 2).

Table 2. Presents the aggregated causal contributions

Policy Construct	Effect (β ± SE)	Significance	Direction
Fiscal Subsidy	+0.47 ± 0.05	$p < 0.01$	Positive
Co-Production	+0.32 ± 0.05	$p < 0.05$	Positive
Infrastructure	+0.29 ± 0.04	$p < 0.05$	Positive
Content Restriction	−0.36 ± 0.04	$p < 0.05$	Negative
Import Control	−0.34 ± 0.06	$p < 0.05$	Negative

Counterfactual simulation showed that replacing "restriction-oriented" language with "innovation-oriented" terms in the 2020 Film Industrial Guidance could increase predicted investment by 14.8 ± 2.1% and production volume by 11.2 ± 1.9% (figure 1). Policy synergy tests revealed multiplicative effects where combined fiscal and innovation signals enhanced output by 18.4 ± 2.5%, exceeding the additive expectation ($p < 0.01$).

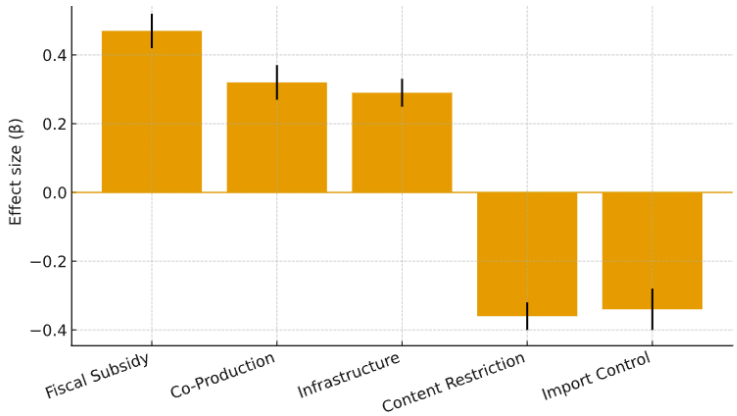

Figure 1. Causal effect estimates of policy constructs on market outcomes (β ± SE)

5.3. Cross-sector and temporal generalization

To test generalization, the CRL model was applied to 500 cultural tourism and creative economy policies. Without retraining, predictive accuracy remained high (MAE = 0.197 ± 0.014; CES = 0.911 ± 0.008), indicating transferability. Temporal analysis showed that causal coupling strength between policy intent and market outcomes peaked during regulatory transitions: 2014–2016 (0.42 ± 0.07) and 2020–2022 (0.45 ± 0.06). These periods coincide with institutional reforms and post-pandemic

DOI: 10.54254/2755-2721/2025.29779

recovery, demonstrating that linguistic framing dynamically interacts with macroeconomic uncertainty.

The stability of causal representation was further verified by temporal cross-validation, showing an $R^2 = 0.89 \pm 0.03$ between predicted and actual outcomes. These findings confirm that semantic-level causal learning captures enduring linguistic mechanisms that extend beyond individual policy contexts.

6. Conclusion

This study has constructed a brand-new causal representation learning framework that can directly infer market responses based on the semantics of policy texts. By integrating transformer-based embeddings, structural causal models, and counterfactual simulations, this framework builds a bridge between computational linguistics and economic modeling. The new approach makes contributions at both the methodological and substantive levels. It demonstrates that causal reasoning can be extended to the textual representation of economic explanations.

It leads to the view of how language can serve as a regulatory tool with measurable economic impacts. Subsequent research can expand this framework to policy comparisons among various countries and add multimodal signals (for example, legal commentary or media sentiment). The research results show that Some language structures in policy rhetoric can have measurable and directional impacts on market activities. Discourse guided by subsidies and innovation can promote a positive investment state, while growth-oriented remarks are constrained by restrictive remarks. It can also optimize real-time policy predictions, combining text-based causal logic with market analysis, which is the key to achieving transparent and evidence-based cultural governance.

Author contribution

Xinye Xu and Xuanqing Huang contributed equally to this paper.

References

[1] Yang, Z., & Shen, B. (2025). Estimating textual treatment effect via causal disentangled representation learning. The Journal of Supercomputing, 81(2), 386.
[2] Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward causal representation learning. Proceedings of the IEEE, 109(5), 612-634.
[3] Imai, K., & Nakamura, K. (2024). Causal representation learning with generative artificial intelligence: Application to texts as treatments. arXiv preprint arXiv: 2410.00903.
[4] Zhou, Y., & He, Y. (2023). Causal inference from text: Unveiling interactions between variables. arXiv preprint arXiv: 2311.05286.
[5] Drury, B., Oliveira, H. G., & de Andrade Lopes, A. (2022). A survey of the extraction and applications of causal relations. Natural Language Engineering, 28(3), 361-400.
[6] Dell, M. (2025). Deep learning for economists. Journal of Economic Literature, 63(1), 5-58.
[7] Chu, Z., & Li, S. (2023). Causal effect estimation: Recent progress, challenges, and opportunities. Machine Learning for Causal Inference, 79-100.
[8] Manzoor, M. A., Albarri, S., Xian, Z., Meng, Z., Nakov, P., & Liang, S. (2023). Multimodality representation learning: A survey on evolution, pretraining and its applications. ACM Transactions on Multimedia Computing, Communications and Applications, 20(3), 1-34.
[9] Sá, S. L. D., Rocha, A. A. D. A., & Paes, A. (2021). Predicting popularity of video streaming services with representation learning: A survey and a real-world case study. Sensors, 21(21), 7328.
[10] Hartmann, J., & Netzer, O. (2023). Natural language processing in marketing. In Artificial intelligence in marketing (pp. 191-215). Emerald Publishing Limited.