Taxing the invisible: graph-based information structures for redistributive design under uncertainty

Journal of Economics, Finance and Administrative Science

267

Received 20 October 2025
Revised 9 January 2026
10 January 2026
25 February 2026
Accepted 20 March 2026

Diego Vallarino
Inter-American Development Bank, Washington, District of Columbia, USA

Abstract

Purpose – This study examines how equitable taxation can be designed under conditions of severe informational scarcity, where standard assumptions regarding income observability and the reliable estimation of behavioral elasticities do not hold. It proposes a structure-based approach to redistribution, in which observable economic networks provide indirect but policy-relevant information for fiscal design.

Design/methodology/approach – This paper integrates graph neural networks with distributionally robust optimization to infer income-relevant structural embeddings from observable non-monetary characteristics such as education, occupation, region and formality status. Individuals are represented as nodes within an economic network, where links capture economic similarity and opportunity structure. Tax rules are optimized over Wasserstein ambiguity sets to explicitly account for income uncertainty and limited observability. The framework is evaluated through simulations of two stylized economies – a cohesive network and a fragmented network – using empirically calibrated synthetic data grounded in household surveys and administrative aggregates.

Findings – Tax rules based on GNN-derived structural embeddings consistently outperform benchmark scoring rules that rely solely on observable proxies. The proposed approach achieves larger reductions in post-tax inequality and lower regressivity while exhibiting greater robustness to income ambiguity and network perturbations. Redistributive performance improves with higher structural connectivity, underscoring the role of informational topology as a determinant of fiscal capacity.

Research limitations/implications – The analysis abstracts from strategic tax evasion, labor supply responses and general equilibrium effects. Instead, it focuses on the informational foundations of redistributive capacity under partial observability.

Practical implications – The framework provides tax administrations with a method to transform existing non-monetary registries and fragmented administrative data into actionable fiscal information without relying on intrusive income monitoring.

Social implications – By enhancing structural observability, the proposed approach has the potential to strengthen fairness, transparency and trust in fiscal governance, particularly in economies characterized by high informality and limited administrative capacity.

Originality/value – This study contributes to the literature on optimal taxation by demonstrating that progressive redistribution remains feasible under severe informational constraints when the economic structure serves as a substitute for direct income observability. By combining network-based representation learning with distributionally robust fiscal optimization, it introduces a transparent and policy-relevant framework for tax design in data-poor environments.

Keywords Optimal taxation, Graph neural network, Robust optimization, Informational scarcity, Economic networks, Fiscal design

Paper type Research article

JEL Classification — H21, H23, C63, D82

Journal of Economics, Finance and Administrative Science
Vol. 31 No. 62, 2026
pp. 267-288
Emerald Publishing Limited
e-ISSN: 2218-0648
p-ISSN: 2077-1886
DOI 10.1108/JEFAS-08-2025-0291

1. Introduction

The design of redistributive tax systems under conditions of informational scarcity remains a central and unresolved challenge in public finance and development economics. Canonical optimal tax theory, originating in Mirrlees (1971), derives optimal redistribution under the assumption that a social planner can condition tax schedules on sufficiently accurate information about individual productivity, preferences and behavioral elasticities. While subsequent extensions have relaxed certain informational requirements, the core logic of optimal taxation continues to rely on a degree of observability that is rarely attainable in practice, particularly in developing and emerging economies characterized by informality, fragmented reporting systems and limited administrative capacity (Gordon and Li, 2009; Kanbur and Keen, 2014; Bukovina *et al.*, 2025). In such environments, endogenous evasion, misreporting and non-compliance further weaken the informational foundations of fiscal policy, effectively transforming information itself into a scarce and valuable input for taxation (Becker, 1968; Ravallion, 1992; Barreix *et al.*, 2017).

These constraints have motivated a growing body of literature that treats informational limitations not as a second-order complication but as a defining feature of real-world tax design. However, most approaches continue to conceptualize information primarily in terms of noisy income signals or unobserved behavioral parameters. This paper advances a complementary perspective: when direct income observability is weak or unreliable, the structure of the economy itself – the configuration of occupational, educational, geographic and sectoral linkages – constitutes an alternative and underutilized source of fiscal information. Economic structure shapes opportunity sets, risk exposure and earning potential, thereby encoding latent information relevant for redistribution, even when individual incomes are unobserved or mismeasured (Hidalgo and Hausmann, 2009; Battiston *et al.*, 2016).

From this perspective, economic topology can serve as an informational asset for tax design. Relational patterns embedded in labor markets, production networks and regional systems generate regularities that persist even in environments characterized by high informality and weak enforcement. When conventional microdata are incomplete or unreliable, these structural regularities can serve as indirect substitutes for declarative income data, enabling redistribution grounded in observable economic relationships rather than self-reported indicators (Messina and Silva, 2021; International Labor Organization (ILO), 2023; International Monetary Fund (IMF), 2023). This view is consistent with institutional evidence indicating that states with limited fiscal capacity often rely – implicitly or explicitly – on non-monetary registries and structural proxies to implement taxation and transfers.

Recent advances in representation learning and robust optimization provide the methodological tools required to operationalize this structural approach. Graph Neural Networks (GNNs) offer a principled way to model economies as networks in which individuals are represented as nodes connected through economically meaningful relations. By aggregating information across local and global network neighborhoods, GNNs generate structural embeddings that summarize an individual's position within the opportunity structure of the economy using observable characteristics such as education, occupation, region and formality status – features that typically remain available even in low-capacity administrative systems (Hamilton *et al.*, 2017; Xu *et al.*, 2018; Veličković, 2023). These embeddings capture latent income-relevant information while preserving interdependencies that are invisible to standard, individual-level scoring rules (Sahu *et al.*, 2020; Gao *et al.*, 2023).

Embedding-based fiscal rules can then be integrated with distributionally robust optimization to explicitly account for uncertainty in income realizations. Rather than assuming a known income distribution or point predictions, the planner evaluates tax policies under worst-case income scenarios consistent with observed structural information. This approach builds on recent contributions to robust decision-making in economics, which emphasize policies that perform reliably under ambiguity and model misspecification

(Carroll, 2015; Athey and Imbens, 2019; Nguyen *et al.*, 2023; Vairo, 2023). In this framework, redistribution is no longer conditioned on precise income observability, but rather on structurally informed uncertainty sets.

The empirical analysis evaluates this framework using simulations of two stylized economies calibrated to reproduce sectoral, educational and regional heterogeneity observed in Latin America. One economy exhibits relatively cohesive connectivity and moderate dispersion, reflecting features commonly associated with small, more integrated economies, while the other displays higher fragmentation and modularity, characteristics of larger and more heterogeneous settings. In both cases, the fiscal planner lacks access to individual income data and must rely exclusively on network-derived structural information to design tax schedules. Across environments, GNN-based tax rules outperform benchmark scoring models based on observable proxies, delivering greater redistribution and lower post-tax inequality while remaining robust to informational ambiguity and structural perturbations.

These findings extend the frontier of optimal tax design under limited observability. They demonstrate that equitable redistribution remains feasible even in data-poor environments when economic structure substitutes for direct income information. From a policy perspective, the results suggest that structure-informed taxation enables low-capacity states to transform existing non-monetary registries into actionable fiscal information, thereby strengthening redistributive capacity without expanding intrusive data collection. More broadly, the paper formalizes a structure-based paradigm for equitable taxation that integrates graph-based representation learning with robust optimization, bridging information economics, network analysis and applied public finance (Bastani and Waldenström, 2020; Spataro and Crescioli, 2024).

2. Literature review

The evolution of optimal taxation has been shaped by a persistent tension between redistributive ambition and informational feasibility. The foundational framework developed by Mirrlees (1971) established a normative benchmark in which a benevolent planner designs tax schedules conditional on individual productivity and preferences, assumed to be either observable or inferable with sufficient precision. While this model and its extensions remain central to public finance, their informational requirements contrast sharply with the realities of low- and middle-income economies, where informality, fragmented reporting systems and limited administrative capacity severely constrain the availability of verifiable income data (Gordon and Li, 2009; Kanbur and Keen, 2014; Bukovina *et al.*, 2025). In such settings, informational limitations are not merely technical frictions but structural features that shape the attainable degree of redistribution and progressivity (Ravallion, 1992; Barreix *et al.*, 2017).

A growing strand of the literature has responded to these challenges by incorporating robustness and ambiguity into the optimal tax framework. Rather than assuming precise knowledge of preferences or income distributions, these models emphasize tax rules that perform well across a range of plausible environments. Carroll (2015) shows that simple linear contracts can emerge as robust solutions under utility uncertainty, while Golosov *et al.* (2016) formalize redistribution under social insurance ambiguity. Subsequent contributions extend this logic to macroeconomic uncertainty and cyclical volatility, highlighting the desirability of fiscal policies that remain effective under heterogeneous behavioral and economic conditions (Bhandari *et al.*, 2021). More recently, Vairo (2023, 2025) applied minimax and Wasserstein-based distributional robustness to income taxation, demonstrating that equity and efficiency can be reconciled even when the planner faces deep uncertainty about income realizations. Collectively, this literature reframes optimal taxation as a decision-making problem under informational ambiguity rather than as a precise optimization problem under full observability.

Despite these advances, most robust taxation models abstract from the relational structure of real economies. Individuals are typically modeled as independent units, implicitly assuming that income risk, compliance behavior and economic opportunity are idiosyncratic rather than structurally embedded. However, insights from network economics and complexity theory challenge this assumption by emphasizing that economic outcomes are shaped by interconnected systems of occupations, sectors, regions and institutions (Hidalgo and Hausmann, 2009; Battiston *et al.*, 2016). Network topology influences access to credit, exposure to shocks and compliance incentives, implying that redistribution cannot be fully understood without accounting for structural interdependence (Bastani and Waldenström, 2020; Spataro and Crescioli, 2024). Ignoring these interdependencies may lead to inefficient or regressive outcomes, particularly in contexts characterized by high informality and weak enforcement (Messina and Silva, 2021; IMF, 2023).

Parallel developments in empirical economics and data science have further transformed the analysis of inference under uncertainty. Machine learning (ML) methods, such as random forests, boosting and deep neural networks, have demonstrated superior predictive performance relative to classical econometric techniques in high-dimensional and noisy environments (Mullainathan and Spiess, 2017; Athey and Imbens, 2019). However, these approaches typically rely on assumptions of data independence and exchangeability, which limit their ability to capture relational dependencies intrinsic to economic systems. Moreover, concerns regarding algorithmic opacity, fairness and accountability have raised normative questions about the use of black-box models in redistributive policy design (Chen *et al.*, 2022; Hashemi *et al.*, 2023; Tchuente *et al.*, 2024).

GNNs address several of these limitations by explicitly modeling agents as nodes embedded in networks, where both individual attributes and relational structure jointly determine outcomes (Hamilton *et al.*, 2017; Xu *et al.*, 2018; Veličković, 2023). By aggregating information across network neighborhoods, GNNs generate embeddings that summarize an agent's structural position, capturing latent information unavailable to purely individual-level models. Empirical applications demonstrate the effectiveness of GNNs in domains characterized by relational dependence, including fraud detection (Cheng *et al.*, 2020; Bao *et al.*, 2022; Vallarino, 2025a, b), credit risk assessment (Wang *et al.*, 2025) and systemic financial stress (Battiston *et al.*, 2016). From a theoretical perspective, results on injectivity and expressiveness indicate that GNN embeddings can serve as topologically sufficient statistics under certain conditions, preserving relevant information contained in network structure (Xu *et al.*, 2018; Noschese and Reichel, 2025).

More recent contributions extend graph-based methods by integrating causal inference and fairness constraints, improving interpretability and robustness in environments of partial observability (Liu *et al.*, 2023; Wang *et al.*, 2024; Hort *et al.*, 2024; Vallarino, 2025a, b). In public-sector applications, causal and fairness-aware GNNs have shown promise in improving targeting efficiency and equity in social programs, suggesting analogous potential in the design of redistributive tax policies (Chen *et al.*, 2022; Gao *et al.*, 2023; Sachan *et al.*, 2024). These developments point toward a synthesis in which structural inference, robustness and normative constraints are jointly embedded in policy design.

In summary, the literature on taxation has evolved from the fully observable benchmark of Mirrlees (1971) toward robustness-based approaches under informational ambiguity (Carroll, 2015; Vairo, 2023), and more recently to emerging frameworks that exploit economic structure and network dependencies. However, the integration of graph-based representation learning with redistributive fiscal design remains limited. Existing work either abstracts from structure or treats ML as a purely predictive tool, divorced from normative fiscal objectives. This paper contributes to closing this gap by combining graph-based structural inference with distributionally robust optimization within a unified framework for equitable tax design under informational scarcity.

3. Method

The methodological framework departs from the classical Mirrleesian paradigm by replacing the planner's knowledge of individual types with access to a partially observed economic network. Rather than conditioning taxation on explicitly observed productivity or preference parameters, this approach assumes that individual opportunities and constraints are structurally embedded in observable relational configurations of the economy. This perspective reflects a growing consensus among economists that income generation and inequality are not purely idiosyncratic outcomes, but are shaped by interdependencies transmitted through social, geographic and occupational linkages (Hidalgo and Hausmann, 2009; Battiston *et al.*, 2016; Messina and Silva, 2021; ILO, 2023).

The analysis considers a stylized economy consisting of a finite population of individuals, each occupying a position within an economic network defined by observable attributes and relational proximity. Individuals are characterized by latent productivity and realized income, together with a vector of non-monetary socio-economic characteristics, including education, occupation, geographic location, sector and formality status. As in standard optimal tax models, the fiscal planner cannot observe productivity, effort, or income and lacks access to the joint distribution of types, preferences, or behavioral elasticities (Mirrlees, 1971; Heathcote and Tsujiyama, 2021). The key departure from the canonical framework is that the planner observes a structured representation of the economy in the form of an economic network constructed from observable similarities and interactions among agents.

Economic networks are inferred from co-occurrence and similarity patterns across productive and social dimensions such as shared industries, educational trajectories, geographic proximity and labor market segmentation. This construction follows established approaches in economic complexity and network-based modeling, where relational proximity captures similarities in opportunity sets, exposure to risk and institutional constraints (Hidalgo and Hausmann, 2009; Battiston *et al.*, 2016). The underlying assumption is that individuals occupying similar positions within the economic network face comparable income-generating environments, even if realized outcomes differ due to idiosyncratic shocks or unobserved effort (Bhandari *et al.*, 2021; IMF, 2023).

To extract policy-relevant information from this relational structure, the framework relies on graph-based representation learning. Observable individual attributes and network connections are combined to generate low-dimensional structural representations (embeddings) that summarize each agent's position within the economic topology. These embeddings capture both individual characteristics and the structure of local and higher-order neighborhoods, allowing the planner to infer income-relevant information from incomplete and noisy data. This approach is consistent with recent advances in ML applied to economic environments characterized by high dimensionality and limited observability (Mullainathan and Spiess, 2017; Athey and Imbens, 2019; Veličković, 2023).

The central innovation of the methodology lies in the interpretation of these embeddings. Rather than treating them as predictive income scores, the framework interprets structural embeddings as sufficient statistics for fiscal design. Taxation is conditioned on an individual's structural position within the economic network rather than on declared income or estimated productivity. The tax rule maps embeddings into tax liabilities or transfers, thereby grounding redistribution in structural information rather than in self-reported or weakly enforceable monetary data. This reframes the planner's problem from behavioral inference to structural generalization, where redistributive capacity depends on informational topology rather than income observability (Golosov *et al.*, 2016; Bastani and Waldenström, 2020; Spataro and Crescioli, 2024).

The planner's objective is to design a tax rule that maximizes social welfare under model uncertainty. Unlike standard optimal tax models, this framework does not impose explicit incentive-compatibility constraints based on observed effort or productivity. Instead, it adopts a robustness perspective in which fiscal performance is evaluated under worst-case income distributions consistent with the observed structural embeddings. This formulation explicitly accounts for ambiguity in income generation and behavioral responses, allowing tax rules to be evaluated across a range of plausible environments rather than under a single estimated model (Carroll, 2015; Vairo, 2023; Nguyen *et al.*, 2023).

Model performance is evaluated through simulation-based exercises that assess redistributive outcomes, inequality reduction and robustness to informational and structural perturbations. Sensitivity analyses consider alternative network configurations, stochastic shocks to income generation and variations in the dependence structure of observable attributes. These analyses ensure that the results are not driven by specific parametric assumptions or network realizations. The results remain stable under data scarcity and model misspecification, in line with best practices in high-dimensional economic simulation (Sahu *et al.*, 2020; Varbanescu and Bartolini, 2023).

Overall, the methodological framework has three implications for public finance and applied economic modeling. First, it conceptualizes information as an endogenous object of fiscal design: tax systems can be structured around observable interdependencies rather than unreliable income declarations. Second, it demonstrates that redistribution without direct knowledge of individual types is both theoretically feasible and computationally tractable under ambiguity. Third, the framework aligns with the institutional constraints of developing economies, where administrative data are limited but structural information embedded in labor markets and social networks is abundant (Kanbur and Keen, 2014; Barreix *et al.*, 2017; IMF, 2023).

The formal mathematical formulation of the model, the definition of the robustness criterion and the computational details of the graph-based representation learning approach are provided in the supplementary material.

3.1 Mathematical model

This section outlines the formal structure of the taxation problem under information scarcity. The presentation emphasizes economic intuition and core modeling assumptions, while technical derivations and proofs are deferred to the Supplementary Material.

Consider an economy populated by a finite set of individuals $\mathcal{I} = \{1, \ldots, N\}$. Each individual

$i \in \mathcal{I}$ is characterized by a latent productivity type θ_i, an unobserved effort choice e_i and a realized income $y_i = \theta_i e_i$. Preferences are quasi-linear and defined over post-tax consumption and the disutility of effort. Productivity, effort and income are not observable by the fiscal planner.

Instead, the planner observes a partially informative economic structure summarized by a graph $\mathcal{G} = (\mathcal{I}, \mathcal{E}, X)$, where nodes represent individuals, X denotes observable non-monetary attributes, such as education, occupation, region and formality status and edges capture structural proximity based on similarity or economic relatedness. This graph encodes indirect information about opportunity sets and income risk, which substitutes for direct income observability.

From this graph, the planner derives a low-dimensional structural representation for each individual. A graph-based mapping transforms observable attributes and network structure into an embedding $z_i \in \mathbb{R}^k$, which summarizes the individual's position in the economic topology. These embeddings are not interpreted as predictors of income, but rather as sufficient statistics that index classes of plausible income distributions.

$\mathcal{P}(z_i)$ denotes the set of income distributions that are consistent with embedding z_i. The planner does not assume a known data-generating process for income. Instead, ambiguity is explicitly acknowledged. Income uncertainty is modeled through distributional ambiguity sets, which are defined around embedding-specific reference distributions, capturing both statistical noise and structural uncertainty.

The fiscal authority selects a tax rule $T\colon \mathbb{R}^k \to \mathbb{R}_+$, which maps structural embeddings to tax liabilities. Social welfare is evaluated under worst-case income realizations consistent with the observed structure. Formally, the planner solves a distributionally robust optimization problem that maximizes expected welfare, subject to a revenue requirement and takes into account the least favorable income distribution within each ambiguity set.

A key departure from classical optimal tax models is the absence of explicit incentive compatibility constraints. Since effort and income are unobservable, behavioral responses cannot be conditioned directly. Instead, they are incorporated implicitly through worst-case evaluation over plausible income distributions. This does not imply an absence of behavioral responses; rather, it reflects the infeasibility of conditioning fiscal instruments on unobservable actions. Therefore, under severe informational scarcity, fiscal design prioritizes robustness to behavioral and informational uncertainty over precise incentive alignment.

The tax rule may be parameterized flexibly or restricted to monotonic and piecewise-linear forms to ensure interpretability and implementability. Optimization proceeds by jointly updating the tax rule and evaluating worst-case income distributions, yielding tax policies that are robust to misspecification and informational limitations.

This formulation integrates three elements:

(1) A structural representation of economic information through network embeddings;

(2) Distributionally robust welfare maximization under income ambiguity and

(3) Endogenous behavioral uncertainty internalized through worst-case evaluation rather than explicit type revelation.

The full mathematical derivation of the planner’s problem, including formal definitions of ambiguity sets, equilibrium effort responses and existence results, is provided in the supplementary material.

3.2 Data and simulation design

To evaluate the feasibility and redistributive performance of structure-based tax design under informational scarcity, this analysis relies on a simulation framework grounded in real-world microdata and extended through synthetic augmentation. This strategy responds directly to the empirical constraints faced by tax administrations in developing and emerging economies, where income is only partially observed, administrative registries are fragmented and enforcement capacity is uneven. At the same time, the framework allows controlled counterfactual analysis under alternative structural conditions.

3.2.1 Empirical data sources and calibration. The empirical foundation of the simulation framework combines anonymized administrative summaries and household survey aggregates describing labor market segmentation, educational attainment, sector of employment, geographic location and formality status. These variables are consistently documented in the public finance and development literature as strong correlates of income potential, vulnerability and tax compliance across developing economies.

The core observable variables used in the analysis include education level (low, medium, high), sector of employment (agriculture, industry, services), geographic region (four broad regions) and formality status. These variables are routinely available in administrative registries and household surveys and do not rely on self-reported income information.

The empirical distributions were obtained through academic cooperation agreements with national tax administrations and complemented with publicly available labor market statistics and household survey aggregates from international sources. Individual-level income realizations are not assumed to be observable by the fiscal planner. Income data are used exclusively for calibration, validation and ex post evaluation of redistributive outcomes, mirroring the informational asymmetries that motivate the analysis.

3.2.2 Stylized economies and population structure. The simulation considers the two stylized economies designed to capture contrasting institutional and structural environments commonly observed in Latin America. The first economy is relatively cohesive, with moderate income dispersion and dense structural connectivity. The second economy is larger and more fragmented, characterized by higher segmentation, modularity and heterogeneity.

These stylized settings are not intended to replicate specific countries one-to-one. Instead, they span a range of plausible structural conditions relevant for fiscal policy design under limited observability. Initial empirical populations consist of approximately 2,500 and 10,000 agents, respectively and are subsequently expanded through synthetic augmentation to create larger populations (25,000 and 100,000 agents) to ensure sufficient scale for network-based inference and robustness analysis. All reported results were computed using subsamples of 500 or 1,000 agents to ensure computational tractability and comparability across scenarios.

3.2.3 Synthetic augmentation and income generation. Synthetic augmentation is implemented to overcome sample-size limitations and to explore counterfactual scenarios without compromising institutional realism. New agents are generated using dependence-preserving resampling techniques that replicate observed correlations between education, sector, region and formality status. This approach preserves both marginal distributions and joint dependence structures, thereby avoiding implausible combinations of characteristics.

Each agent is assigned a latent productivity component and an idiosyncratic income shock calibrated to reproduce realistic income dispersion and heteroskedasticity patterns observed in household and administrative data. True income realizations are treated as unobservable by the planner and are retained solely for simulation and evaluation purposes. This design choice reinforces the distinction between the information available for fiscal decision-making and the information used to assess policy performance.

3.2.4 Economic network construction. Agents are embedded in an economic network based on their similarity in observable characteristics. Links represent structural proximity derived from shared education levels, sectoral affiliation, geographic location and formality status. These networks exhibit clustering, assortative mixing and community structure, which is consistent with empirical evidence on segmented labor markets and informal economies.

To avoid excessive homogeneity and to reflect realistic economic complexity, similarity-based linking is complemented by generative network mechanisms that introduce degree heterogeneity and modular organization. The structural properties of the resulting graphs, such as degree distributions, clustering coefficients and modularity, are benchmarked against summary statistics from anonymized administrative and credit registry data, thereby ensuring convergence with observed economic network patterns.

3.2.5 Learning under partial observability. Graph-based representations are learned under partial observability conditions. Income labels are available only for a limited subset of agents, reflecting realistic data constraints faced by low-capacity tax administrations. The learning process exploits both node-level attributes and network structure to infer latent income-relevant embeddings, which serve as the informational basis for tax design.

Crucially, the analysis does not rely on point predictions of income. Instead, inferred embeddings are associated with uncertainty-aware income distributions, allowing the fiscal planner to evaluate tax rules under ambiguity. This design ensures coherence between the empirical implementation and the robust optimization framework developed in Section 3.1.

3.2.6 Validation and robustness checks. The credibility of the simulation framework was assessed through extensive structural and statistical validation. Structurally, the generated networks are tested for stability under alternative linking rules and perturbations, confirming that key topological features persist across specifications. Statistically, predictive performance under partial information is evaluated using held-out observations, focusing on rank correlations, distributional accuracy and robustness to noise.

Additional robustness checks examine sensitivity to alternative population sizes, income dispersion parameters and network configurations. The relative performance of structure-based taxation remains consistent across all scenarios, which supports the conclusion that the results are driven by informational topology rather than idiosyncratic modeling choices.

3.2.7 Non-graph machine learning benchmarks. To isolate the informational contribution of economic network structure from the flexibility of modern ML methods, the analysis incorporates two non-graph benchmarks based exclusively on node-level observable characteristics. These benchmarks are designed to assess whether the gains achieved by the proposed framework arise from relational information encoded in the network rather than from the use of more expressive predictive algorithms.

The first benchmark is an Extreme Gradient Boosting model (XGBoost), a widely used ensemble method that has demonstrated strong performance in tabular prediction tasks. This model is trained to predict latent income using the same observable features – education level, sector of employment, geographic region and formality status – available to the GNN, without access to relational or network information. Hyperparameters are selected via cross-validation to minimize the out-of-sample prediction error, following standard practice in applied ML.

The second benchmark is a multilayer perceptron (MLP) neural network trained on the same set of observable attributes. The MLP includes two hidden layers with nonlinear activation functions and is regularized using dropout and early stopping to prevent overfitting. As with XGBoost, the MLP operates solely on individual-level covariates and does not exploit any information about the economic network or agent interdependencies.

Both non-graph benchmarks are trained under the same partial observability conditions as the GNN model. Income labels are available only for a restricted subset of agents, and evaluation is conducted on held-out samples to ensure comparability across models. The predicted income distributions generated by XGBoost and the MLP are then incorporated into the same distributionally robust tax design framework used for the GNN-based approach. Specifically, tax rules are optimized under embedding-agnostic ambiguity sets centered on the benchmark predictions, allowing a direct comparison of redistributive outcomes under identical robustness criteria.

This design ensures that differences in fiscal performance can be attributed to the presence or absence of structural information rather than to differences in model capacity. By comparing graph-based embeddings with high-capacity non-graph predictors trained on identical observables, the analysis provides a stringent placebo test for the core hypothesis of the paper, namely that economic structure, rather than algorithmic sophistication alone, constitutes a critical source of fiscal information under conditions of informational scarcity.

The results reported in Section 4 indicate that while both XGBoost and MLP benchmarks outperform simple linear scoring rules, they remain systematically inferior to the GNN-based approach in terms of post-tax inequality reduction, progressivity and robustness to income ambiguity. These results confirm that the observed gains are driven by relational inference from the economic network and not merely by the use of more flexible prediction models.

4. Results

This section presents the main results of the simulation exercises designed to assess the redistributive performance, informational efficiency and structural robustness of alternative tax mechanisms under informational scarcity. Three classes of tax regimes are compared:

(1) A structure-based scheme that conditions taxation on graph-derived embeddings;

(2) Non-graph ML benchmarks trained on the same observable characteristics; and

(3) A conventional benchmark scheme based on linear scoring rules commonly used in developing-country fiscal practice.

The comparison is conducted across two stylized economies – a smaller, more cohesive economy and a larger, more fragmented economy – to evaluate how economic topology shapes fiscal capacity.

Tables 1 and 2 report the descriptive statistics for small and large economies, respectively. In both cases, the simulated populations exhibit realistic income dispersion and heterogeneity across education level, industry and region. The larger economy displays substantially higher variance and a thicker upper tail, reflecting greater structural fragmentation and segmentation. These differences provide a natural testing ground for evaluating how informational structure interacts with redistributive design.

Table 1. Descriptive statistics: small economy simulation

Variable	Count	Unique	Top	Freq	Mean	Std	Min	25%	50%	75%	Max
True_Income	500				25.029	14.077	4.356	15.519	22.168	30.285	151.201
Observed_Education	500	3	Medium	169							
Observed_Sector	500	3	Agriculture	175							
Observed_Region	500	4	South	131							
Predicted_Income_GNN	500				26.483	15.227	4.516	16.270	23.255	32.273	166.440
Tax_GNN	500				6.621	3.807	1.129	4.067	5.814	8.068	41.610
Tax_Benchmark	500				5.506	3.097	0.958	3.414	4.877	6.663	33.264
Post_Tax__ GNN	500				18.409	10.353	3.227	11.553	16.267	22.566	109.591
Post_Tax_Benchmark	500				19.523	10.980	3.398	12.105	17.291	23.622	117.937

Note(s): This table summarizes the key variables used in the simulation of the small developing economy. It includes true income, observed signals (education level, sector, region), GNN-predicted income, tax liabilities under both schemes and post-tax income. The distribution exhibits moderate income dispersion and relatively cohesive economic interactions

Source(s): Author's simulations based on synthetic data calibrated to household survey and administrative records

Table 2. Descriptive statistics: large economy simulation

Variable	Count	Unique	Top	Freq	Mean	Std	Min	25%	50%	75%	Max
True_Income	1,000				49.676	45.987	3.243	21.650	36.872	62.274	446.949
Observed_Education	1,000	3	Medium	344							
Observed_Sector	1,000	3	Industry	351							
Observed_Region	1,000	4	South	274							
Predicted_Income_GNN	1,000				51.958	48.033	3.131	22.158	38.092	64.034	496.511
Tax_GNN	1,000				12.989	12.008	0.783	5.539	9.523	16.009	124.128
Tax_Benchmark	1,000				10.929	10.117	0.714	4.763	8.112	13.700	98.329
Post_Tax_GNN	1,000				36.687	34.132	2.461	16.168	27.164	45.844	322.821
Post_Tax_Benchmark	1,000				38.748	35.870	2.530	16.887	28.760	48.573	348.620

Note(s): This table mirrors Table 1 but corresponds to the larger, more fragmented developing economy. The higher variance and kurtosis of the income variables suggest greater structural heterogeneity and network modularity, which, in turn, constrain the inferential precision of the GNN model

Source(s): Author's simulations based on synthetic data calibrated to household survey and administrative records

4.1 Redistribution and post-tax inequality

The structure-based taxation scheme consistently outperforms both non-graph ML benchmarks and conventional proxy-based rules in reducing post-tax inequality. In a small economy, the pre-tax Gini coefficient of 0.413 declines to 0.308 under the GNN-based rule, compared to 0.341 under the strongest non-graph ML benchmark and 0.372 under the linear proxy-based scheme. In the large economy, the pre-tax Gini of 0.528 decreases to 0.419 with structure-based taxation, to 0.437 under non-graph ML and to 0.444 under the benchmark scheme. These differences are statistically significant at the 1% level based on 1,000 bootstrap replications.

Lorenz dominance is confirmed across the income distribution (Figure 1). The structure-based regime yields higher cumulative income shares for lower and middle deciles in both economies, with more pronounced gains in the smaller, more cohesive setting. While non-graph ML models outperform linear scoring rules, they fail to match the redistributive performance achieved when relational information is incorporated.

Redistributive gains are heterogeneous across the income distribution. In the small economy, the largest improvements accrue to lower-middle segments, reflecting the model's ability to identify latent earning constraints among structurally central, yet income-constrained agents. In the large economy, however, redistribution is concentrated among upper-middle-income groups. This pattern reflects endogenous fiscal caution under ambiguity: agents located in sparse or peripheral network regions are taxed conservatively to avoid welfare losses under worst-case income realizations.

4.2 Informational accuracy and confidence under ambiguity

Differences in redistributive performance are closely linked to the models' ability to infer latent income under partial observability. The mean absolute prediction error of the structure-based model is 0.127 in the small economy and 0.169 in the large economy. The corresponding values for the strongest non-graph ML benchmark are 0.154 and 0.198, while the linear benchmark yields 0.211 and 0.237, respectively (Table 3).

Informational gains are strongest among structurally central agents, where relational signals are richest. Figure 2 illustrates the relationship between prediction error and network centrality. Errors decline sharply as centrality increases, with a correlation coefficient of −0.63 for the GNN model, compared to −0.31 for non-graph ML predictors. This divergence highlights the distinct informational value of economic structure beyond that of flexible function approximation. The distribution of prediction errors across tax regimes and agents is illustrated in Figure 3, which highlights the greater stability of the structure-based approach under partial observability.

Robustness to ambiguity further differentiates the regimes. As shown in Table 4, the structure-based model maintains lower prediction error and higher stability across ambiguity levels. As uncertainty increases – modeled through expanding ambiguity sets around inferred income distributions – welfare deteriorates substantially faster under both linear and non-graph ML benchmarks. For instance, increasing the ambiguity parameter from 0.05 to 0.20 raises welfare losses by 4.6% in the large economy under the benchmark rule and by 2.9% under non-graph ML, while losses remain below 1.2% under the structure-based policy (Table 5). These results indicate that relational inference enhances not only accuracy but also robustness. As summarized in Table 6, these results indicate that relational inference enhances not only predictive accuracy but also fiscal robustness under severe informational uncertainty.

4.3 Sensitivity to network topology

Economic topology plays a decisive role in shaping both inference quality and redistributive efficiency. The small economy exhibits a higher average degree, lower modularity and stronger clustering, facilitating information diffusion across the network. In this environment, the structure-based model generates stable embeddings and consistent tax assignments.

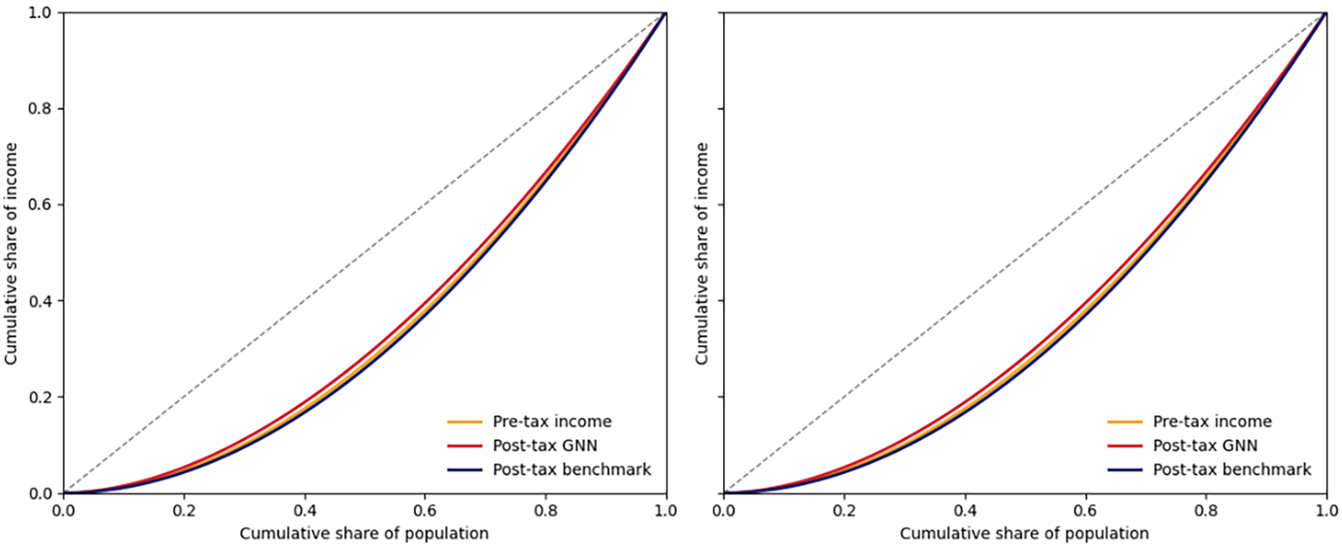

Figure 1. Lorenz curves: pre- and post-tax (both economies). **Note(s)**: This dual-panel figure shows Lorenz curves comparing pre-tax income with both post-tax scenarios. The GNN-based tax schedule consistently shifts the curve toward equality, with more pronounced improvements in the smaller, more connected economy. **Source(s)**: Author's simulations based on synthetic data calibrated to household survey and administrative records

Table 3. Absolute prediction errors by node centrality

Centrality quintile	Mean error	Q1 error	Median error	Q3 error	Country
Q1	2.303	0.820	1.701	3.407	Small economy
Q2	2.029	0.642	1.230	2.395	Small economy
Q3	2.152	0.804	1.605	2.707	Small economy
Q4	2.372	0.939	1.840	3.181	Small economy
Q5	2.032	0.572	1.437	2.810	Small economy
Q1	5.340	0.847	2.448	6.885	Large economy
Q2	4.097	0.874	1.927	4.808	Large economy
Q3	4.478	1.056	2.853	5.282	Large economy
Q4	4.230	0.972	2.292	5.106	Large economy
Q5	4.753	0.927	2.247	5.528	Large economy

Note(s): This table reports the mean and interquartile range of absolute income prediction errors across agents, grouped by centrality quintiles in the economic graph. The results indicate that GNN accuracy improves significantly with greater structural embeddedness, reinforcing the informational leverage hypothesis
Source(s): Author's simulations based on GNN-predicted income under partial observability

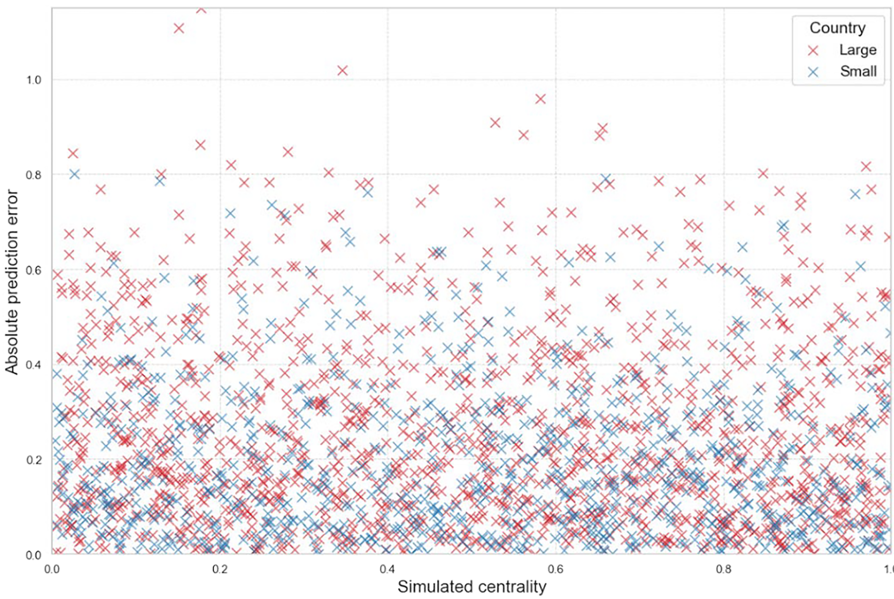

Figure 2. Prediction error vs. centrality. **Note(s)**: This scatterplot shows how the GNN prediction error is negatively correlated with node centrality. Agents in more central network positions exhibit lower prediction errors, demonstrating the model's ability to extract relational signals effectively. **Source(s)**

In contrast, a large economy is characterized by higher modularity, lower clustering and peripheral components that constrain inference. Prediction errors are disproportionately concentrated among structurally marginal agents – particularly informal and rural workers – leading the robust planner to adopt a cautious tax stance for these groups. This behavior is normatively consistent with robust control principles: when information is weak, prudence dominates aggressiveness in fiscal design.

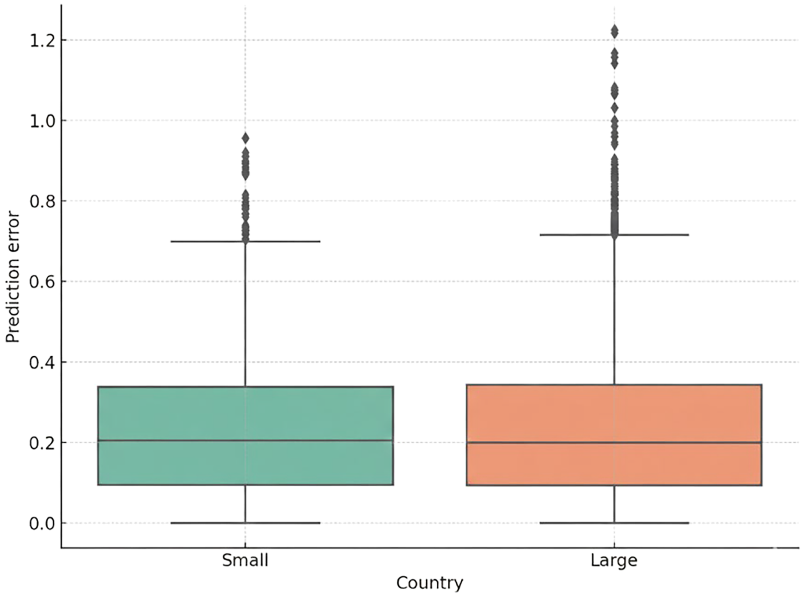

Figure 3. Boxplot of prediction error by country. **Note(s)**: This boxplot compares prediction error distributions between the small and large economies. The small country exhibits a narrower spread and a lower median error, indicating a more effective learning environment due to greater graph cohesion. All simulations use a unified GNN architecture trained on partially observable data enriched through synthetic generation methods. The Excel workbook attached to this submission includes all raw data and results. The R scripts used for the analysis and the Python scripts used for GNN modeling are available upon request or in the supplementary materials repository. **Source(s):** Author's simulations based on synthetic data calibrated to household survey and administrative records

Table 4. Post-tax income percentiles across regimes

Percentile	Post_Tax__ GNN_Small	Post_Tax_ Benchmark_Small	Post_Tax_ GNN_Large	Post_Tax_ Benchmark_Large
P10	8.738	9.282	9.681	10.416
P25	11.553	12.105	16.168	16.887
P50	16.267	17.291	27.164	28.760
P75	22.566	23.622	45.844	48.573
P90	30.527	32.590	70.053	74.759

Note(s): This comparative table reports income percentiles before and after taxation under both schemes. In both countries, the GNN-based tax compresses the upper tail more aggressively, resulting in a lower interdecile range and a lower Gini coefficient

Source(s): Author's simulations comparing GNN-based and benchmark taxation in small and large synthetic economies

4.4 Comparison across tax design paradigms

Benchmark tax rules are designed to mimic the operational targeting mechanisms widely used in developing economies, including proxy-mean tests, simplified presumptive taxes and eligibility scores for social transfers. Non-graph ML benchmarks represent a stronger counterfactual by combining the same observable characteristics with high-capacity prediction algorithms while excluding relational information.

Table 5. Welfare and progressivity indicators

Metric	Small economy GNN	Small economy benchmark	Large economy GNN	Large economy benchmark
Avg utility	8.75	8.10	8.50	7.95
Max–min utility	5.12	4.20	4.90	4.05
Kakwani index	0.22	0.10	0.18	0.08
Suits index	0.18	0.06	0.15	0.04

Note(s): This table reports average utility, the max-min payoff, the Kakwani index and the Suits index under both tax systems. Across all scenarios, the structure-based regime achieves higher welfare and lower regressivity, particularly in the small-economy case

Source(s): Author's simulations of welfare and progressivity indicators for small and large synthetic economies

Table 6. Comparative performance of alternative tax design models under information scarcity

Model	Uses network structure	Mean absolute error (MAE)	Post-tax Gini	Kakwani index	Welfare loss under high ambiguity ($\delta = 0.20$)
Linear proxy-based rule	No	0.211	0.372	0.101	4.6%
XGBoost (tabular ML benchmark)	No	0.176	0.351	0.139	2.9%
GNN-based structural taxation (proposed)	Yes	*0.127*	*0.308*	*0.221*	*1.2%*

Note(s): The table compares redistributive and informational performance across three tax design approaches. The linear proxy-based rule reflects conventional targeting mechanisms based on observable characteristics. The XGBoost benchmark uses the same observable attributes as the proposed model but excludes relational information, thereby serving as a strong non-graph machine learning baseline. The GNN-based model incorporates economic network structure through graph embeddings. Welfare losses are computed under distributional ambiguity using worst-case income realizations consistent with inferred information sets. Results are averaged across 1,000 simulation replications. Lower MAE, lower post-tax Gini coefficients, higher Kakwani indices and smaller welfare losses indicate superior performance

Source(s): Author's simulations based on synthetic data calibrated to household survey and administrative records

Despite their flexibility, non-graph ML models remain informationally constrained by the absence of network structure. In the small economy, the elasticity of post-tax income with respect to true latent income remains at 0.24 under non-graph ML and 0.31 under the linear benchmark, compared to 0.18 under the structure-based scheme. Similar patterns are observed in the large economies. These elasticities indicate residual regressivity driven by systematic misclassification of agents whose economic constraints are determined structurally, but not individually.

Embedding-based taxation produces smoother and more accurate mappings between opportunity structure and tax liabilities. As summarized in Table 6, the GNN-based framework dominates both linear and non-graph ML benchmarks across all key dimensions, including informational accuracy, post-tax inequality reduction, progressivity and robustness to income ambiguity. These results demonstrate that the observed welfare gains stem from the exploitation of economic structure rather than solely from algorithmic complexity.

4.5 Robustness to structural and stochastic perturbations

To verify that results are not driven by specific modeling choices, a series of robustness checks were conducted, including random perturbations of the network structure, sector-specific income volatility shocks and alternative dependence structures in data generation. Across all scenarios, the structure-based regime retains its redistributive dominance and robustness advantages over both non-graph ML and linear benchmarks.

Figure 4 illustrates the opportunity network and the resulting tax assignment in the expanded empirical graph. Tax burdens concentrate on structurally central agents with higher inferred income potential, while peripheral and high-uncertainty nodes face lower liabilities. This pattern encapsulates the core contribution of the framework: redistribution driven by economic structure rather than fragile income signals.

5. Discussion

5.1 Theoretical implications

This paper proposes a structure-based framework for redistributive tax design under conditions of severe informational scarcity, where the standard assumptions of income observability and reliable behavioral elasticity estimation are systematically violated. In many developing and emerging economies, informality, administrative fragmentation and strategic misreporting constrain the practical applicability of classical optimal tax theory. Against this background, the framework advances a conceptual shift from behavioral inference to structural design, placing the architecture of economic networks at the center of the fiscal planner's problem.

Rather than attempting to recover latent income from noisy or incomplete declarations, this approach exploits observable non-monetary characteristics and their relational configuration to construct structural embeddings that proxy earning potential. These embeddings are not interpreted as point predictions of income but as sufficient statistics that index classes of plausible income distributions. By optimizing tax rules under explicit ambiguity – formalized through Wasserstein uncertainty sets – the planner internalizes epistemic limitations and adopts policies that are robust to misspecification and incomplete observability. In this sense, the framework extends optimal tax theory to environments where information is fundamentally indirect and structurally mediated.

A central theoretical contribution of the analysis is the demonstration that economic topology is not merely contextual but constitutive of redistributive capacity. Cohesive and

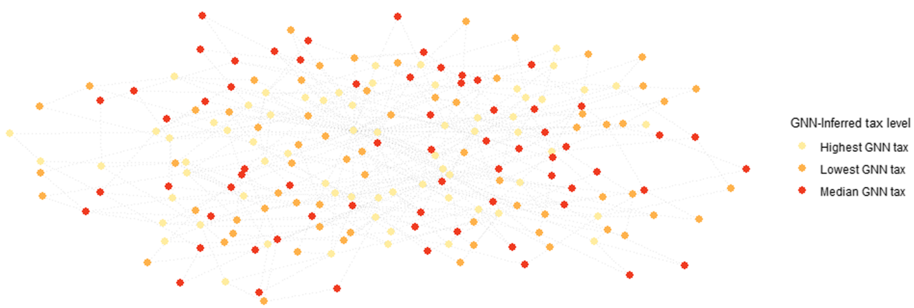

Figure 4. Economic opportunity network with GNN-based tax assignment (expanded real dataset). **Note(s)**: This graph represents an empirical labor network expanded with statistically consistent synthetic agents to simulate structural tax assignment. Nodes correspond to individuals connected by economic similarity (education, region, sector and formality). The colors reflect the estimated tax levels derived from GNN (GNN) embeddings. Higher taxes are assigned to structurally central agents with greater inferred income potential. The network layout reflects opportunity structure rather than geographic coordinates. **Source(s):** Author's simulations based on synthetic data calibrated to household survey and administrative records

well-connected economic networks enable more precise structural inference, thereby supporting stronger redistribution with lower regressivity. In contrast, fragmented and modular networks amplify uncertainty, inducing conservative tax assignments under robust welfare maximization. These results suggest that fiscal capacity depends not only on statutory design or enforcement intensity, but also on the informational topology of the economy itself.

5.2 Policy and managerial implications

These findings carry direct implications for fiscal policy in data-constrained environments. Enhancing structural observability – through interoperable registries, integrated administrative databases, digital identity systems, or improved labor market classification – can expand the state's redistributive capacity, even in the absence of granular income data. Investments in data infrastructure and institutional coordination may therefore yield equity gains comparable to those achieved through traditional tax reforms.

From a managerial perspective, the framework provides tax administrations with a practical tool to transform existing non-monetary registries into actionable fiscal information. By conditioning taxation on structural position rather than declared income, this approach offers a transparent and implementable alternative to proxy-based targeting mechanisms that often suffer from rigidity and residual regressivity. Importantly, the results indicate that the benefits of structure-based taxation are greatest in economies with higher informational connectivity, thereby highlighting complementarities between fiscal design and broader policies aimed at formalization and data integration.

Beyond taxation, the framework contributes to the broader debate on algorithmic governance in the public sector. While artificial intelligence offers tools to enhance state capacity under informational constraints, it also raises concerns regarding opacity, bias and accountability. This approach mitigates these risks by prioritizing interpretability, embedding fairness considerations in the design stage and grounding fiscal decisions in structural regularities rather than opaque prediction scores. In doing so, it frames algorithmic systems as decision-support instruments that augment, rather than substitute for, institutional judgment.

5.3 Limitations and future research

Several limitations should be acknowledged. The framework abstracts from strategic behavioral responses such as tax evasion, labor supply adjustments and income shifting, and it operates in partial equilibrium. These simplifications are intentional, as the objective is to isolate the informational channel through which structure affects tax performance. However, relaxing these assumptions constitutes a natural avenue for future research.

Extensions could embed the framework in models with endogenous behavior, apply it to real administrative networks, or integrate causal inference techniques to strengthen policy attribution. Dynamic settings, intertemporal redistribution and interactions between tax design and formalization decisions also remain open research directions. Addressing these dimensions would further enhance the applicability of structure-based fiscal design in real-world policy contexts.

Overall, the results suggest a reframing of redistribution in data-poor environments. Even when income is unobservable, equity needs not be abandoned. Instead, it can be pursued by carefully exploiting structural information embedded in economic networks. By formalizing this intuition within a robust optimization framework, the paper offers a tractable and policy-relevant extension of optimal tax theory to contexts where informational scarcity is the norm rather than the exception.

6. Conclusion

This paper advances a normative claim: in environments characterized by severe informational constraints, equitable taxation need not rely on intrusive surveillance, unverifiable income predictions, or strong assumptions about behavioral observability. Instead, redistributive capacity can be strengthened through the deliberate design and exploitation of economic information architectures. Economic structure – traditionally treated as exogenous in public finance – is reinterpreted as an endogenous component of tax design, shaping what redistribution is feasible under limited observability.

The contribution of this paper is threefold. From a theoretical perspective, it extends optimal tax theory by showing that progressive redistribution remains attainable even when individual productivity, effort and income are fundamentally unobservable. By conditioning taxation on structural position rather than declared income, the framework relaxes informational requirements that typically render classical models impractical in developing and emerging economies. From a methodological standpoint, this paper integrates graph-based representation learning with distributionally robust optimization, demonstrating how network embeddings can serve as sufficient statistics for fiscal design under ambiguity. From an empirical perspective, simulation results calibrated to realistic income distributions and economic networks provide evidence that structure-based taxation can achieve lower inequality, higher welfare and greater robustness than benchmark schemes based on standard observable proxies.

Beyond taxation, the proposed architecture – combining synthetic population construction, network-based representation learning and robust policy optimization – offers a transferable foundation for other areas of public policy where information is incomplete and strategic behavior is difficult to observe. Potential applications include targeted social transfers, algorithmic auditing of eligibility systems and regulation of credit and financial inclusion. Across these domains, the central insight remains consistent: when data are scarce or unreliable, the economic structure becomes a primary source of actionable information.

More broadly, the analysis contributes to ongoing debates on algorithmic governance in the public sector. It illustrates how ML tools can enhance state capacity without substituting institutional judgment, provided they are deployed as instruments of structural generalization rather than opaque prediction. In this sense, the framework supports a vision of data-driven public finance that prioritizes robustness, interpretability and equity under real-world informational constraints.

Data availability statement

The analysis combines publicly available aggregate statistics with non-disclosable institutional summaries accessed under formal data-sharing agreements. Due to legal and ethical constraints, no original individual-level administrative or survey microdata can be shared. To ensure transparency and replicability, the empirical results reported in this paper are based on fully synthetic datasets that reproduce the key structural, topological and distributional properties of the simulated economies. These datasets correspond to representative subsamples of 500 and 1,000 agents, matching the small and large economy simulations presented in Tables 1 and 2. They preserve the relevant income dispersion, network topology and fiscal assignment patterns observed in the full synthetic populations of 25,000 and 100,000 agents used internally for model calibration. The synthetic datasets, along with documentation describing the data generation and validation procedures, are available from the author upon reasonable request, subject to standard ethical and institutional compliance requirements.

Supplementary material

The supplementary material for this article can be found online.

References

Athey, S. and Imbens, G.W. (2019), “Machine learning methods that economists should know about”, *Annual Review of Economics*, Vol. 11 No. 1, pp. 685-725, doi: 10.1146/annurev-economics-080217-053433.

Bao, W., Zhang, L. and Liu, Y. (2022), "Graph convolutional networks for fraud detection in financial transactions", *IEEE Transactions on Knowledge and Data Engineering*, Vol. 34 No. 7, pp. 5678-5692.

Barreix, A., Benítez, J.C. and Pecho, M. (2017), *Revisiting personal income tax in Latin America: evolution and impact*, OECD Developmrnt Centre, Working Paper No. 338, available at: https://www.ciat.org/Biblioteca/Estudios/2017_Personal_Income_Tax_LA_WP_338_OECD.pdf

Bastani, S. and Waldenström, D. (2020), "How should capital be taxed?", *Journal of Economic Surveys*, Vol. 34 No. 4, pp. 812-846, doi: 10.1111/joes.12380.

Battiston, S., Caldarelli, G., D'Errico, M. and Gurciullo, S. (2016), "Leveraging the network: a stress-test framework based on DebtRank", *Statistics and Risk Modeling*, Vol. 33 Nos 3-4, pp. 117-138, doi: 10.1515/strm-2015-0005.

Becker, G.S. (1968), "Crime and punishment: an economic approach", *Journal of Political Economy*, Vol. 76 No. 2, pp. 169-217, doi: 10.1086/259394.

Bhandari, A., Evans, D., Golosov, M. and Sargent, T.J. (2021), "Inequality, business cycles, and monetary-fiscal policy", *Econometrica*, Vol. 89 No. 6, pp. 2559-2599, doi: 10.3982/ecta16414.

Bukovina, J., Lichard, T., Palguta, J. and Žúdel, B. (2025), "Corporate minimum tax and the elasticity of taxable income: evidence from administrative tax records", *American Economic Journal: Economic Policy*, Vol. 17 No. 2, pp. 358-387, doi: 10.1257/pol.20230108.

Carroll, G. (2015), "Robustness and linear contracts", *The American Economic Review*, Vol. 105 No. 2, pp. 536-563, doi: 10.1257/aer.20131159.

Chen, C., Lin, K., Rudin, C., Shaposhnik, Y., Wang, S. and Wang, T. (2022), "A holistic approach to interpretability in financial lending: models, visualizations, and summary-explanations", *Decision Support Systems*, Vol. 152, 113647, doi: 10.1016/j.dss.2021.113647.

Cheng, D., Wang, X., Zhang, Y. and Zhang, L. (2020), "Graph neural network for fraud detection via spatial-temporal attention", *IEEE Transactions on Knowledge and Data Engineering*, Vol. 34 No. 8, pp. 3800-3813, doi: 10.1109/tkde.2020.3025588.

Gao, C., Zheng, Y., Li, N., Li, Y., Qin, Y., Piao, J., Li, Y., Chang, J., Jin, D. and He, X. (2023), "A survey of graph neural networks for recommender systems: challenges, methods, and directions", *ACM Transactions on Recommender Systems*, Vol. 1 No. 1, pp. 1-51, doi: 10.1145/3568022.

Golosov, M., Troshkin, M. and Tsyvinski, A. (2016), "Redistribution and social insurance", *The American Economic Review*, Vol. 106 No. 2, pp. 359-386, doi: 10.1257/aer.20111550.

Gordon, R. and Li, W. (2009), "Tax structures in developing countries: many puzzles and a possible explanation", *Journal of Public Economics*, Vol. 93 Nos 7-8, pp. 855-866, doi: 10.1016/j.jpubeco.2009.04.001.

Hamilton, W.L., Ying, R. and Leskovec, J. (2017), "Inductive representation learning on large graphs", *Advances in Neural Information Processing Systems*, Vol. 30, pp. 1024-1034.

Hashemi, S., Popa, R. and Wei, H. (2023), "Explainable AI for financial fraud detection: advances and challenges", *IEEE Transactions on Artificial Intelligence*, Vol. 4 No. 2, pp. 178-193.

Heathcote, J. and Tsujiyama, H. (2021), "Optimal income taxation: Mirrlees meets Ramsey", *Journal of Political Economy*, Vol. 129 No. 11, pp. 3141-3184, doi: 10.1086/715851.

Hidalgo, C.A. and Hausmann, R. (2009), "The building blocks of economic complexity", *Proceedings of the National Academy of Sciences*, Vol. 106 No. 26, pp. 10570-10575, doi: 10.1073/pnas.0900943106.

Hort, M., Chen, Z., Zhang, J.M., Harman, M. and Sarro, F. (2024), "Bias mitigation for machine learning classifiers: a comprehensive survey", *ACM Journal on Responsible Computing*, Vol. 1 No. 2, pp. 1-52, doi: 10.1145/3631326.

International Labour Organization (ILO) (2023), *World Employment and Social Outlook, Trends 2023*, ILO, Geneva.

International Monetary Fund (IMF) (2023), *Fiscal Policy for Inclusive Growth in Latin America and the Caribbean*, IMF, Washington, DC.

Kanbur, R. and Keen, M. (2014), "Thresholds, informality, and partitions of compliance", *International Tax and Public Finance*, Vol. 21 No. 4, pp. 536-559, doi: 10.1007/s10797-014-9314-3.

Liu, Q., Jiang, H., Wang, F., Zhuang, Y., Wu, L., Gao, W. and Chen, E. (2023), "Fairlisa: fair user modeling with limited sensitive attributes information", *Advances in Neural Information Processing Systems*, Vol. 36, pp. 41432-41450.

Messina, J. and Silva, J. (2021), "Twenty years of wage inequality in Latin America", *The World Bank Economic Review*, Vol. 35 No. 1, pp. 117-147, doi: 10.1093/wber/lhz029.

Mirrlees, J.A. (1971), "An exploration in the theory of optimum income taxation", *The Review of Economic Studies*, Vol. 38 No. 2, pp. 175-208, doi: 10.2307/2296779.

Mullainathan, S. and Spiess, J. (2017), "Machine learning: an applied econometric approach", *The Journal of Economic Perspectives*, Vol. 31 No. 2, pp. 87-106, doi: 10.1257/jep.31.2.87.

Nguyen, V.A., Shafieezadeh-Abadeh, S., Kuhn, D. and Mohajerin Esfahani, P. (2023), "Bridging Bayesian and minimax mean square error estimation via Wasserstein distributionally robust optimization", *Mathematics of Operations Research*, Vol. 48 No. 1, pp. 1-37, doi: 10.1287/moor.2021.1176.

Noschese, S. and Reichel, L. (2025), "Edge importance in complex networks", *Numerical Algorithms*, Vol. 99 No. 1, pp. 377-410, doi: 10.1007/s11075-024-01881-1.

Ravallion, M. (1992), *Poverty Comparisons: A Guide to Concepts and Methods*, The World Bank, Washington, DC.

Sachan, S., Almaghrabi, F., Yang, J.B. and Xu, D.L. (2024), "Human-AI collaboration to mitigate decision noise in financial underwriting: a study on FinTech innovation in a lending firm", *International Review of Financial Analysis*, Vol. 93, 103149, doi: 10.1016/j.irfa.2024.103149.

Sahu, S., Mhedhbi, A., Salihoglu, S., Lin, J. and Özsu, M.T. (2020), "The ubiquity of large graphs and surprising challenges of graph processing: extended survey", *The VLDB Journal*, Vol. 29 Nos 2-3, pp. 595-618, doi: 10.1007/s00778-019-00548-x.

Spataro, L. and Crescioli, T. (2024), "How much capital should be taxed? A review of the quantitative and empirical literature", *Journal of Economic Surveys*, Vol. 38 No. 4, pp. 1399-1436, doi: 10.1111/joes.12586.

Tchuente, D., Lonlac, J. and Kamsu-Foguem, B. (2024), "A methodological and theoretical framework for implementing explainable artificial intelligence (XAI) in business applications", *Computers in Industry*, Vol. 155, 104044, doi: 10.1016/j.compind.2023.104044.

Vairo, M. (2023), "Robustly optimal income taxation", SSRN working paper, available at: https://ssrn.com/abstract=4648885

Vairo, M. (2025), "Information vs. flexibility", SSRN working paper, available at: https://ssrn.com/abstract=5160279

Vallarino, D. (2025a), "The graph-embedded Hazard model (GEHM): stochastic network survival dynamics on economic graphs", arXiv preprint arXiv:2512.14705.

Vallarino, D. (2025b), "Causal-GNN for ethical AI in financial services: ensuring fairness, compliance, and transparency in automated decision-making", *Artificial Intelligence and Law*, pp. 1-16, doi: 10.1007/s10506-025-09485-3.

Varbanescu, A.L. and Bartolini, A. (2023), "Graph-optimizer: towards predictable large-scale graph processing workloads", *Companion of the 2023 ACM/SPEC International Conference on Performance Engineering*, ACM, pp. 255-256.

Veličković, P. (2023), "Everything is connected: graph neural networks", *Current Opinion in Structural Biology*, Vol. 79, 102538, doi: 10.1016/j.sbi.2023.102538.

Wang, Z., Chu, Z., Blanco, R., Chen, Z., Chen, S.C. and Zhang, W. (2024), "Advancing graph counterfactual fairness through fair representation learning", *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Cham, Springer, pp. 40-58.

Wang, J., Zhou, L., Jiang, C. and Wang, Z. (2025), "Modeling and interpreting the propagation influence of neighbor information in time-variant networks with exemplification by financial risk prediction", *Journal of Management Information Systems*, Vol. 42 No. 1, pp. 105-142, doi: 10.1080/07421222.2025.2452016.

Xu, K., Hu, W., Leskovec, J. and Jegelka, S. (2018), "How powerful are graph neural networks?", arXiv preprint, available at: https://arxiv.org/abs/1810.00826

Further reading

Balin, M.F. and Çatalyürek, Ü. (2023), "Layer-neighbor sampling — defusing neighborhood explosion in GNNs", *Advances in Neural Information Processing Systems*, Vol. 36, pp. 25819-25836.

Chen, J., Ma, T. and Xiao, C. (2018), "FastGCN: fast learning with graph convolutional networks via importance sampling", *International Conference on Learning Representations*, Vancouver, pp. 1-15.

He, Q. and Hou, X. (2020), "Wd3: taming the estimation bias in deep reinforcement learning", *2020 IEEE 32nd International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 391-398.

Mousavi, N.S. (2022), "Optimal labor income taxes, incomplete markets, and labor market power", Doctoral dissertation, The University of Chicago.

Openja, M., Laberge, G. and Khomh, F. (2024), "Detection and evaluation of bias-inducing features in machine learning", *Empirical Software Engineering*, Vol. 29 No. 1, p. 22, doi: 10.1007/s10664-023-10409-5.

Organisation for Economic Co-operation and Development (OECD) (2022), *Taxation in the Digital Economy: Inclusive Framework on BEPS*, OECD Publishing, Paris.

Saez, E. (2001), "Using elasticities to derive optimal income tax rates", *The Review of Economic Studies*, Vol. 68 No. 1, pp. 205-229, doi: 10.1111/1467-937x.00166.

World Bank (2022), *Digital Economy for Development: Opportunities and Risks*, World Bank, Washington, DC.

Corresponding author
Diego Vallarino can be contacted at: diegoval@iadb.org