



Ca' Foscari
University
of Venice


**Department
of Economics**

Working Paper

**Adriano Amati
Monica Billio
Marco Di Cataldo
Giovanni Mastrobuoni**

**Mafia Connections:
Infiltration in Corporate
Ownership**

ISSN: 1827-3580
No. 24/WP/2026





Mafia Connections: Infiltration in Corporate Ownership

Adriano Amati

ETH Zürich

Monica Billio

Ca' Foscari University of Venice

Marco Di Cataldo

Ca' Foscari University of Venice; London School of Economics

Giovanni Mastrobuoni

Collegio Carlo Alberto

Abstract

Learned representations have transformed the measurement of unstructured data in economics. We extend it to relational data, showing that Temporal Graph Networks encode economically meaningful behavior in dynamic corporate ownership networks. Using a high-resolution Italian ownership graph anchored on 5,700 judicially confiscated firms, we train a TGN that is never supervised on confiscation to produce time-varying firm embeddings, and we summarize their geometry with an Infiltration Proximity Index (IPI): a real-time measure of how densely a firm's latent neighborhood is populated by firms whose confiscation is already legally known. We validate the index along three dimensions. It forecasts confiscation out of sample up to four years ahead, with a higher area under the ROC curve than the full set of firm-level financial variables at every horizon and for every classifier, and with far fewer missed confiscations at the cost of flagging more firms that are never confiscated; the geometry it summarizes places firms confiscated only later closer to firms already confiscated at the time of measurement; and it responds coherently to local changes in ownership. We then use the index to date firms' transitions into a high-risk regime in a staggered difference-in-differences design. Around the dated transition, firms display sharp scale expansion, rising liabilities and receivables, cost reallocation, and persistent illiquidity, with only temporary profit gains, patterns consistent with firms operating as conduits for financial flows rather than as profit maximizers.

Keywords

corporate network, organized crime, infiltration, graph neural networks, embeddings

JEL Codes

C45, C55, D85, L1, K42, L25, G32

Address for correspondence:

Marco Di Cataldo

Department of Economics
Ca' Foscari University of Venice
Cannaregio 873, Fondamenta S. Giobbe
30121 Venezia - Italy
e-mail: marco.dicataldo@unive.it

This Working Paper is published under the auspices of the Department of Economics of the Ca' Foscari University of Venice. Opinions expressed herein are those of the authors and not those of the Department. The Working Paper series is designed to divulge preliminary or incomplete work, circulated to favour discussion and comments. Citation of this paper should consider its provisional character.

Mafia Connections: Infiltration in Corporate Ownership

Adriano Amati¹, Monica Billio², Marco Di Cataldo^{2,3}, Giovanni Mastrobuoni⁴

¹*ETH Zürich*

²*Ca' Foscari University of Venice*

³*London School of Economics*

⁴*Collegio Carlo Alberto*

Abstract

Learned representations have transformed the measurement of unstructured data in economics. We extend it to relational data, showing that Temporal Graph Networks encode economically meaningful behavior in dynamic corporate ownership networks. Using a high-resolution Italian ownership graph anchored on 5,700 judicially confiscated firms, we train a TGN that is never supervised on confiscation to produce time-varying firm embeddings, and we summarize their geometry with an Infiltration Proximity Index (IPI): a real-time measure of how densely a firm's latent neighborhood is populated by firms whose confiscation is already legally known. We validate the index along three dimensions. It forecasts confiscation out of sample up to four years ahead, with a higher area under the ROC curve than the full set of firm-level financial variables at every horizon and for every classifier, and with far fewer missed confiscations at the cost of flagging more firms that are never confiscated; the geometry it summarizes places firms confiscated only later closer to firms already confiscated at the time of measurement; and it responds coherently to local changes in ownership. We then use the index to date firms' transitions into a high-risk regime in a staggered difference-in-differences design. Around the dated transition, firms display sharp scale expansion, rising liabilities and receivables, cost reallocation, and persistent illiquidity, with only temporary profit gains, patterns consistent with firms operating as conduits for financial flows rather than as profit maximizers.

JEL Classification: C45, C55, D85, L1, K42, L25, G32

Keywords: corporate network, organized crime, infiltration, graph neural networks, embeddings.¹

Email addresses: adriano.amati@ethz.ch (Adriano Amati), billio@unive.it (Monica Billio), marco.dicataldo@unive.it (Marco Di Cataldo), giovanni.mastrobuoni@carloalberto.org (Giovanni Mastrobuoni)

¹*Acknowledgments:* We thank Mario Padula, Paolo Buonanno, Ruben Durante, Paolo Pinotti, Rocco Macchiavello and all participants to presentations at Ca' Foscari, Napoli Parthenope, Verona, Padova, Sofia, Palermo for useful comments and suggestions. We thank Dr Gaetano Merenda and Dr Stefano Capone (ANBSC) for sharing data on confiscation timelines. We are grateful to Infocamere and Unioncamere and in particular Dr Niccolo' Stamboglis for sharing data on confiscated firms' network. We also thank Guardia di Finanza Regione Veneto for insights and collaboration. All errors are our own.

1. Introduction

In recent years, the economics literature has increasingly adopted machine learning methods—particularly neural networks—to enrich both empirical and theoretical analysis (Dell, 2025; Zheng et al., 2023; Ash and Hansen, 2023; Gentzkow et al., 2019a). These tools have been widely used to improve forecasting accuracy and to transform previously unstructured data, such as text, audio, and images, into measurable economic indicators through the use of embeddings. Applications range from the use of textual embeddings to analyze political and institutional language (Gentzkow et al., 2019b; Ash et al., 2024), to the use of audio features to capture information beyond the text in central bank communications (Gorodnichenko et al., 2023), and to the application of embedding techniques to economic measurement problems such as demand estimation (Magnolfi et al., 2025). Despite this rapid expansion, the use of graph representation learning—and especially graph neural network (GNN) embeddings—remains remarkably underexplored in economics. While such methods have shown strong potential for uncovering latent structures and network effects (Brogaard and Chen, 2024; Villarraga and Daziano, 2025), their integration into economic research is still at an early stage, offering a promising yet largely untapped avenue for measurement and inference in networked economic systems.

We take up this challenge by focusing on a classic example of such a networked system: corporate ownership. Our core idea is that the complex structure of these networks—where nodes represent firms or individuals and edges relationships through equity stakes—can itself be encoded and studied through learned representations. GNNs are ideally suited to transform these high-dimensional and evolving ownership structures into firm embeddings. These are compact, learned representations that summarize a firm’s structural position, history of connections, and exposure to other entities in the economy, capturing both local and global dependencies. This approach allows the model to uncover hidden features—such as indirect ownership chains, recurring control patterns, or rapid ownership changes, that are otherwise hard to track directly. We apply this framework to the study of organized crime groups (OCGs) infiltration in the Italian economy, using infiltration as a case study to show how learned ownership embeddings can uncover subtle and dynamic network patterns associated with illicit control and firm vulnerability.

The paper proceeds in four steps. We first construct representations of the dynamic ownership network that are never supervised on confiscation and summarize their geometry in a single firm-level index; we then validate that index’s geometric, predictive, and relational content; we use it to date firm-specific transitions into a high-risk regime; and we finally characterize how firm outcomes evolve around those transitions. Correspondingly, the paper makes four main contributions. First, we construct a novel and uniquely detailed dataset that has not been analyzed before. We build a *dynamic ownership network* for Italy that traces both firm–firm and individual–firm links across more than two decades. Crucially, our network is anchored on judicially confirmed confiscation events, ensuring a ground truth free of false positives. From these confiscated “seed” firms, we reconstruct the surrounding corporate ecosystem by linking all direct and second-order owners and subsidiaries—encompassing both *Persone Giuridiche* (firms) and *Persone Fisiche* (individuals). The resulting network is directed, weighted, and time-stamped at the daily level, yielding a high-resolution view of corporate control and its evolution over time. By integrating this structure with firm-level financial data, we obtain a rich, attributed graph that allows us to study how the dynamics of ownership connectivity themselves can serve as early indicators of infiltration and subsequent confiscation.

Second, we introduce a representation-based measure of latent infiltration risk that exploits the dynamic structure of corporate ownership, and we subject it to a sequence of validation exercises. Building on recent advances in temporal graph learning (Rossi et al., 2020), we use a GNN to learn low-dimensional ownership embeddings for each firm—statistical representations that summarize its evolving position within the ownership network. Importantly, the model is never supervised on confiscation: its training objectives concern the dynamics of ownership relations—the formation of links, the revision of ownership stakes, and their dissolution—so the representations encode structural and temporal regularities of corporate control without being fitted to infiltration labels. We train two variants of the model that differ only in how much firm-level information they see. In the *full-data* variant, every year’s balance sheet and the firm’s current judicial status enter as attribute updates, so the representation summarizes both the ownership and the attribute history of a firm; in the *first-year* variant, a firm’s attributes are observed only once, in its first year, and no judicial information ever enters, so the representation subsequently moves with ownership events alone. We summarize the geometry of the learned space with an Infiltration Proximity Index (IPI): a contin-

uous, real-time measure of how densely a firm’s latent neighborhood is populated by firms whose confiscation is already legally known at the time of measurement. We show that the geometry it summarizes places firms that are confiscated only later closer to the region occupied by firms already confiscated at the time of measurement; that firms’ latent positions move coherently when they first connect to a counterparty that is legally unmarked at the time but confiscated subsequently; and that, when the neighborhood of a confiscated firm subsequently reconnects to previously external firms, the entrants reconnecting to the largest share of that neighborhood occupy distinctive positions in the latent space. These exercises are descriptive validation analyses: they establish what the representations encode, and we attach no causal interpretation to them.

Third, we show that this single index carries forward-looking information about confiscation. In a chronological out-of-sample exercise we forecast whether a firm will be confiscated one to four years ahead, comparing classifiers that use the full set of firm-level characteristics and financial indicators, with five lags of every variable, against classifiers that use the IPI and four of its lags alone. The index attains a higher area under the ROC curve than the firm-level block at every horizon and for every classifier, linear or non-linear. The two information sets also fail in different ways. Firm-level data, when paired with non-linear classifiers, are comparatively precise but miss most of the firms that are later confiscated; the index recovers more than three-quarters of them at every horizon, at the cost of flagging more firms that are never confiscated. In a setting where enforcement is delayed and incomplete, and the underlying infiltration is never observed, this profile—few false negatives, more false positives—is the desirable one: firms flagged by the index but never confiscated display relational behavior consistent with infiltration that was never prosecuted, and a screening device should surface them rather than suppress them.

While this forecasting exercise is conducted in a network constructed around judicially confirmed confiscation cases, the methodological implication is broader. Deploying this representation-learning framework at scale—on the full universe of corporate ownership relations—would provide a systematic screening device for identifying firms whose structural trajectories resemble those of previously infiltrated entities. Such a tool is not intended to substitute investigative or judicial activity, but to complement it: it can serve as a quantitative support instrument to help authorities prioritize monitoring efforts and allocate limited investigative resources toward firms exhibiting elevated structural risk.

Fourth, we use the validated geometry to date the transition into a high-risk regime and deploy this timing variable in a staggered difference-in-differences design. Our objective is not to provide causal estimates of the impact of infiltration, but to document which observable economic characteristics move together with a firm’s first large and persistent increase in embedding-based risk. Focusing on firms that are eventually confiscated, we exploit variation in the timing of IPI jumps to examine how firm-level performance indicators evolve around these transitions. We find that these transitions coincide with rapid expansion in firm scale, but not with sustained improvements in economic fundamentals. Revenues and operating activity grow sharply, profitability rises only temporarily, and liquidity steadily deteriorates. Balance-sheet growth is driven primarily by receivables and short-term liabilities rather than cash or productive assets. This configuration suggests that firms whose embeddings move into high-risk regions of the latent space do not expand in order to accumulate capital or strengthen financial positions. Instead, their balance sheets and operating activity resemble those of high-throughput entities, in which financial and real flows pass through the firm rather than being retained.

These findings align with prior evidence that organized crime infiltration is associated with firm expansion ([Mirenda et al., 2022](#); [Arellano-Bover et al., 2024](#)), but they clarify the economic configuration that accompanies transitions into high-risk structural positions. The combination of rapid scale growth, persistent illiquidity, increasing claims and liabilities, and declining cash ratios points to firms operating as conduits for moving resources across a broader network. Rather than maximizing long-run profitability or accumulating liquid buffers, these firms appear to intermediate transactions, warehouse receivables and obligations, and facilitate the circulation of funds within criminal ecosystems. In this sense, embeddings serve a diagnostic function: they allow us to identify when firms begin to look structurally risky and to trace which economic indicators co-move with this transition into a conduit-like operating regime. More broadly, this approach demonstrates how learned representations of ownership structure can be integrated with standard econometric tools to study the economic manifestation of latent network phenomena that only become legally observable *ex post*.

2. GNN embeddings as a tool to assess organized crime infiltration

Economic research increasingly leverages representation learning to transform unstructured inputs into measurable objects for inference. In macro-finance and monetary economics, for example, acoustic features extracted from central-bank speeches—beyond words alone—predict policy surprises and market moves, showing that audio embeddings contain information orthogonal to text (Gorodnichenko et al., 2023). In political economy, computational-linguistic methods applied to decision-making transcripts recover latent deliberation and persuasion mechanisms at the FOMC, illustrating how high-dimensional text can be mapped into economically interpretable factors (Hansen et al., 2018). Beyond language, recent work embeds product spaces from triplet comparisons and uses those embeddings inside demand systems, demonstrating clear gains in fit and substitution patterns relative to hand-crafted characteristics (Magnolfi et al., 2025). Relatedly, machine-learning pipelines deployed in applied policy contexts (e.g., anti-corruption monitoring) show that learned representations can deliver actionable ex-ante risk scores within standard econometric evaluation frameworks (Ash et al., 2025). Building on this agenda, we introduce graph neural network embeddings of corporate ownership—an unstructured, relational data source—to construct economically meaningful firm-level representations that evolve in continuous time and can be used for prediction and empirical analysis.

We apply this methodological framework to study a phenomenon that critically operates through corporate ownership structures: organized crime (OCG) infiltration. Over the past few decades, OCGs have pushed well beyond their traditional spheres, with a recent Europol report underscoring that 86% of the EU’s most dangerous criminal networks infiltrate lawful enterprises (Europol, 2024). This has drawn significant attention from economists, who are examining how mafia infiltration alters economic incentives and distorts the broader environment (Pinotti, 2015; De Feo and De Luca, 2017; Di Cataldo and Mastrorocco, 2022; Fenizia and Saggio, 2024). A rich literature using administrative and balance-sheet data has begun to identify the empirical signatures of this activity: Mirenda et al. (2022) show that infiltrated firms display revenue surges disconnected from real productivity, followed by financial deterioration, while Le Moglie and Sorrenti

(2022), [Pinotti et al. \(2025\)](#), and [Castelluccio and Rizzica \(2026\)](#) find that infiltration is more likely among smaller, financially fragile firms. Other work frames this as a strategic investment problem, with OCGs choosing how to engage with legal firms based on size and expected returns ([Arellano-Bover et al., 2024](#)).

A critical insight from this literature is that infiltration is not an isolated firm-level event but a network phenomenon. Complementary studies emphasize this relational nature, showing how anti-mafia interventions propagate through firm networks ([Calamunci and Drago, 2020](#)) and how ownership restructuring during crises persists in mafia-prone areas, suggesting infiltration diffuses through these relational structures ([De Luca et al., 2025](#)). This empirical evidence highlights why economic theory has increasingly turned to social network analysis as a powerful framework for understanding and combating crime ([Lindquist and Zenou, 2019](#); [Ballester et al., 2004](#)). Since crime is fundamentally a social phenomenon, the very architecture of a criminal network is theoretically important, as it can either amplify or dampen criminal activity. From a statistical modeling perspective, other research frames this structure as the outcome of a fundamental efficiency-security trade-off. This view posits that organizations must create enough connections to operate efficiently while remaining sparse enough to ensure security and avoid detection ([Lu et al., 2025](#)). This tension shapes their architecture, leading to unique features such as groups of redundant criminals and a mix of overt and deliberately covert ties. Our study builds on this foundation by shifting the focus from the internal social dynamics of criminal-only networks to the economic interactions of OCGs within the legal economy. We consider a mixed environment where criminal actors and legitimate firms interact, allowing us to identify the patterns that signal a firm’s risk of confiscation—whether it was originally created by an OCG or was a legitimate enterprise later preyed upon.

This focus on patterns addresses a fundamental challenge in the study of organized crime infiltration: the infiltration itself is latent and never directly observed. Empirical strategies in the existing literature have relied on partial or proxy indicators—such as ownership transfers involving suspicious individuals ([Mirenda et al., 2022](#)) or persons of interest already identified by authorities ([Cariello et al., 2024](#)), or more structural approaches modelling OCG investment decisions ([Arellano-Bover et al., 2024](#)). While these approaches provide valuable insights, they remain limited to observable manifestations of infiltration ex post. By contrast, our framework infers infiltration risk from the latent dynamics of the ownership network. Using temporal graph networks ([Rossi et al., 2020](#)),

we learn firm-level embeddings that summarize the evolution of ownership relations over time, allowing us to capture subtle and unobserved structural regularities—such as recurrent patterns of control, indirect links, and bursts of ownership turnover—that may precede formal detection or confiscation. This approach enables a shift from static, attribute-based classification toward a dynamic, representation-learning perspective on latent criminal infiltration.

Our study thus bridges (i) AI-based measurement of unstructured economic data, (ii) network analysis of firm behavior, and (iii) the economics of organized crime, contributing to each of them the four elements set out in the Introduction: a novel large-scale dynamic ownership network rooted in judicially confirmed confiscations; a representation-learning method previously unexplored in economics; a measure of latent infiltration risk derived from the geometry of the learned representations and validated along several dimensions; and the use of that measure as a timing device for studying firm outcomes. Together, these move beyond static proxies to capture the dynamic, relational patterns associated with illicit economic activity.

3. Data

This section will briefly describe all the data sources and how they have been used for the study.

Confiscated Firms Data Data on confiscated firms has been provided by the National Agency for Seized and Confiscated Mafia Assets (*Agenzia Nazionale per i beni sequestrati e confiscati alla criminalità organizzata*, ANBSC) which reports detailed data on the universe of firms ever confiscated for illicit activities in relation to OCG. This dataset comprises over 5,700 firms confiscated between 1983 to 2024 and includes key variables such as the percentage of confiscated assets, the date of confiscation, the economic sector in which the firm operates, and precise geo-referenced data, including the firm’s address. Furthermore, the data documents the complete judicial history of the confiscation process, from the initial seizure of assets by authorities (*sequestro*), to the final legal outcome, be it the confirmation of the confiscation through multiple judicial instances or its eventual revocation. Crucially, in contrast to prior research (e.g., [Chiricop et al., 2023](#); [Boeri et al., 2024](#)), our dataset provides an unprecedentedly complete

picture of this phenomenon at a national scale, with the number of confiscated firms being an order of magnitude larger.

For the remainder of this paper, the confiscation date refers to the initial judicial seizure of proprietary assets by authorities, which marks the beginning of legal proceedings against a firm suspected of mafia involvement. Figure A.1 in the Appendix displays the number of confiscations over time. We can observe an increase in confiscations in the years 2009 to 2017 and a sharp decline in later years. This is probably due to the lack of processing of more recent events by ANBSC. Figure A.2 displays the spatial distribution of confiscated firms at the municipality-level activity. We can immediately see that there is a high concentration of confiscated entities in Rome and Milan, followed by Naples and Palermo. Overall, there seems to be a higher concentration of confiscated firms in the South of Italy, which is consistent with past literature on mafia presence in Italy (Arellano-Bover et al., 2024; Pinotti, 2015).

Ownership Network We also collected firm-level data from *InfoCamere*, the Italian company responsible for managing the official business registry. To analyze the economic environment of the confiscated firms, we reconstructed their ownership networks. The network nodes represent either firms (*Persone Giuridiche*, PG) or individuals (*Persone Fisiche*, PF), and the directed edges represent ownership stakes, timestamped at the day level. We mapped the network surrounding each confiscated firm up to two degrees of separation considering all ownership stakes $> 0.05\%$ to define a connection. The reconstruction followed a two-step process:

1. **First-Degree Connections:** For each confiscated firm, we first identified its immediate connections. We compiled a list of all its direct shareholders (those who owned it) and all of its direct subsidiaries (those it owned) over a 20-year period, covering the 10 years before and after the confiscation.
2. **Second-Degree Connections:** For each upstream shareholder, we included all other firms in their portfolio, ensuring these connections fell within the 20-year time window of the original confiscated firm. Similarly, for each downstream subsidiary, we included all of its owners and any other firms it controlled, again respecting the same temporal window. In cases where a second-degree entity was connected to multiple confiscated firms, we retained only the link associated with the earliest confiscation date to provide a unique temporal anchor for our analysis.

This methodology captures a detailed view of the local ownership structure, revealing both the investors and assets linked to the confiscated entities. The resulting network is a weighted—i.e., ownership percentage—, directed—i.e., ownership direction—heterogeneous²—dynamic graph, spanning from 1993 to 2024.³

This network-based approach represents a fundamentally different way to measure the degree of OCG involvement in the legal economy. Rather than attempting to identify potentially infiltrated firms from the entire population (Mirenda et al., 2022; Arellano-Bover et al., 2024), our methodology begins with a solid core of firms whose involvement in criminal activity is judicially confirmed through confiscation. From this reliable anchor of definitively infiltrated entities, we then expand outward, building a targeted network of all connected firms and individuals up to two degrees of separation. The resulting dataset provides a unique view into the immediate economic environment surrounding confirmed OCG activity, enabling us to analyze the firms within this high-risk ecosystem to find predictive signals of infiltration.

Firm-Level Data From *InfoCamere*, we also collect comprehensive details on company registrations, firm characteristics—i.e. number of employees, municipality, industry, and so forth—economic and financial performance as well as balance-sheet data at the yearly level. The two main sources are the business registry (*registro*), with quarterly data covering years from 2001 to 2024, and firm-level company financials (*bilancio*), with yearly data covering the period 2005 to 2024. By integrating the ANBSC confiscation records with the Infocamere ownership and firm-level datasets, we construct a rich, attributed network of yearly snapshots. Figure A.3 displays the distribution across sectors of our sample of firms. We can observe that most confiscated entities operate in construction, real estate, wholesale and manufacturing, with fewer confiscations occurring in the remaining sectors. Table A.1 in the Appendix provides a detailed list of all firm-level variables used in our analysis, categorized into firm characteristics, financial indicators, balance sheet assets, balance sheet liabilities and equity, and income statement items.

The Network By integrating the confiscation records, ownership data, and firm-level

²as it includes both people and firms. In this context people can only be *source* nodes, i.e. they cannot be owned by any other node

³Ideally, this network would be benchmarked against a random sample of the Italian firm population ownership network. However, access to the comprehensive ownership data required to construct such a baseline is presently restricted.

Table 1: Descriptive statistics - Directed Mafia Network

Metric	1993	1998	2003	2008	2013	2015	2017	2019	2021	2023	Whole
Network characteristics:											
Nodes	7,957	11,384	17,819	25,411	25,577	25,366	24,726	23,097	21,340	19,218	36,323
Links	326	3,860	9,271	18,341	13,485	12,686	11,967	9,159	6,988	4,742	40,425
Average Degree	0.041	0.339	0.520	0.722	0.527	0.500	0.484	0.397	0.327	0.247	1.113
Companies (G nodes)	7,757	9,978	14,057	18,189	19,370	19,218	18,881	18,262	17,394	16,197	25,121
Persons (F nodes)	200	1,406	3,762	7,222	6,207	6,148	5,845	4,835	3,946	3,021	11,202
Prop. Companies	0.975	0.876	0.789	0.716	0.757	0.758	0.764	0.791	0.815	0.843	0.692
Confiscated Companies	3,353	3,712	4,380	5,346	5,914	5,957	5,867	5,791	5,676	5,416	6,284
Prop. Confiscated	0.432	0.372	0.312	0.294	0.305	0.310	0.311	0.317	0.326	0.334	0.250
Avg. Confiscation Quota	92.167	91.823	90.909	90.093	90.173	90.246	90.316	90.473	90.566	90.677	90.357
N. of components	7,672 (7631)	8,502 (7517)	10,254 (8544)	10,000 (7779)	13,696 (12074)	14,045 (12694)	14,047 (12762)	14,749 (13930)	14,903 (14355)	14,780 (14472)	4,513 (4576)
Average Clustering Coefficient	0.000 (0.000)	0.002 (0.000)	0.002 (0.000)	0.004 (0.000)	0.003 (0.000)	0.002 (0.000)	0.002 (0.000)	0.002 (0.000)	0.001 (0.000)	0.001 (0.000)	0.009 (0.000)
Overall Clustering Coefficient	0.014 (0.000)	0.003 (0.000)	0.003 (0.000)	0.003 (0.000)	0.005 (0.000)	0.004 (0.000)	0.005 (0.000)	0.005 (0.000)	0.004 (0.000)	0.005 (0.000)	0.004 (0.000)
Largest Component characteristics:											
Nodes	18	123	825	4,472	2,615	4,100	5,610	3,340	2,732	2,267	22,071
Links	26	385	1,388	5,728	3,070	4,708	6,442	3,759	3,018	2,449	28,948
Average Degree	1.444	3.130	1.682	1.281	1.174	1.148	1.148	1.125	1.105	1.080	1.312
Companies (G nodes)	13	106	654	2,811	1,486	1,946	2,772	1,334	958	609	15,222
Persons (F nodes)	5	17	171	1,661	1,129	2,154	2,838	2,006	1,774	1,658	6,849
Prop. Companies	0.722	0.862	0.793	0.629	0.568	0.475	0.494	0.399	0.351	0.269	0.690
Confiscated Companies	10	1	42	166	77	87	146	113	124	110	1,013
Prop. Confiscated	0.769	0.009	0.064	0.059	0.052	0.045	0.053	0.085	0.129	0.181	0.067
Fraction of Nodes	0.002 (0.001)	0.011 (0.003)	0.046 (0.068)	0.176 (0.544)	0.102 (0.092)	0.162 (0.030)	0.227 (0.019)	0.145 (0.003)	0.128 (0.002)	0.118 (0.001)	0.608 (0.849)
Average Clustering Coefficient	0.046 (0.000)	0.000 (0.000)	0.011 (0.000)	0.008 (0.000)	0.007 (0.000)	0.006 (0.000)	0.006 (0.000)	0.004 (0.000)	0.002 (0.000)	0.002 (0.000)	0.011 (0.000)
Overall Clustering Coefficient	0.100 (0.000)	0.000 (0.000)	0.001 (0.000)	0.003 (0.000)	0.006 (0.000)	0.006 (0.000)	0.005 (0.000)	0.005 (0.000)	0.004 (0.000)	0.004 (0.000)	0.004 (0.000)

Notes: The table reports network characteristics for the mafia network for selected years. Number in parenthesis are from simulated Erdős-Renyi random networks with the same number of nodes, links and density as the observed network. Time intervals have been chosen to be every 5 years for older snapshots and every two years for more recent ones. The last column refers to the whole network, i.e. the one including all links and nodes from 1993 to 2023.

financials, we construct the final, rich, attributed network used in our analysis. To explore this complex dataset, we developed an interactive 3D rendering of the graph. The rendering displays the complete static network, which aggregates all connections across the entire period. The static version of the graph counts 40,425 edges and 36,323 nodes, of which 11,202 PFs and 25,121 PGs. Figure A.4a displays the static version of the graph, Figure A.4b displays the static graph’s main component and Table 1 displays


the static and yearly networks’ descriptive statistics. In both Figures A.4a and A.4b, as well as in Table 1, we can observe that there is a giant component accounting for 60% of the nodes and 70% of the edges in the static version of the network. Interestingly, this feature disappears when we look at the yearly snapshots, where the largest component accounts for only 10-15% of the nodes and edges. This stark difference between the long-term aggregate and short-term snapshots may hint at strategic infiltration patterns, such as the use of temporary or rapidly changing ownership structures to evade detection. While a full analysis of this specific network feature is beyond the scope of this paper, it represents a compelling area for future research.

Figures 1 and A.5 provide some visualizations of the network’s yearly snapshots. Figure 1 illustrates the network’s structure in 2016, which is dominated by three prominent hubs. These hubs are all part of a single, large, weakly connected component that, as detailed in the Appendix (Figure A.5d), encompasses nearly 6,000 nodes, just under 7,000 edges, and over 150 confiscated firms. A closer look at the hubs reveals a striking contrast among their most central nodes: two are confiscated micro-firms—i.e. less than 10 employees (indicated by red nodes), while the third is anchored by a large, non-confiscated financial institution (a green node). Panels A.5a to A.5c display other recurrent network structures. Panel A.5a shows a star-like structure centered on a confiscated firm, with many direct connections to both individuals and other firms. Here edges are red to indicate that information on the ownership percentage is missing. Panel A.5b shows a more complex structure, with multiple interconnected firms and individuals, suggesting a web of relationships that could facilitate illicit activities. Panel A.5c displays a dense cluster of non-confiscated firms with a person (white node) at the center, indicating a potentially influential individual within this network.

4. Methodology

This section describes the methodology used to generate firm-level embeddings from the dynamic ownership network. The objective of this approach is to obtain low-dimensional representations that summarize each firm’s evolving position and behavior within the network, while preserving the temporal and structural complexity of ownership relations. To do so, we employ a Temporal Graph Network (TGN) (Rossi et al., 2020), a neural architecture designed for dynamic graphs in which both the links and

Figure 1: Graph Snapshot for 2016



Notes: The figure shows the rendered network for 2016. The legend is displayed on the top right. The 3 most central nodes have been highlighted. Note that each of these hub nodes belong to the same weakly connected component, see [A.5d](#)

node attributes evolve over time. TGNs are particularly well-suited to this setting, as they process ownership changes in continuous time and use the exact timestamps of these events to model how information, control, and influence propagate through the network.

Our main objective is to build a TGN architecture that can learn *ownership embeddings*—low-dimensional vectors that encode each firm’s structural role, relational history, and exposure to other entities in the ownership network. For the model to learn meaningful representations, we supervise it on the three fundamental behaviors that fully characterize edge dynamics in ownership data: (i) *link creation* (a new ownership relation appears), (ii) *link update* (an existing ownership share changes), and (iii) *link removal* (an ownership relation disappears). Each observed event at time t is a tuple $(u, v, t, \text{type}, y_t)$, involving a source node u , i.e. investing firm or person, a target firm node d , the event type $\in \{\text{creation, update, deletion}\}$ and when applicable the ownership stake or its change from the most recent decision.

This design choice is central to our approach. A model trained to predict confiscation would internalize the statistical regularities of law enforcement and judicial detection—learning *where confiscations occur* rather than *how ownership structures evolve*.

Supervising instead on link creation, update, and removal forces the network to infer the mechanisms that govern *who* invests in *whom*, *under what temporal and structural conditions*, and *how* these relationships intensify or dissolve over time—dynamics that precede and often underlie legal outcomes such as confiscation. The resulting embeddings encode each entity’s evolving relational identity: how it participates in the flow of ownership, accumulates or disperses control, and interacts with other actors in its economic environment.

Data Processing The dataset combines two sources of temporal information: (i) daily-stamped ownership relations between firms (Persone Giuridiche, PG) and individuals (Persone Fisiche, PF), and (ii) yearly financial statements for firms. To make these heterogeneous data compatible with a temporal graph model, we convert them into a single chronologically ordered stream of timestamped events.

We first define the universe of nodes by including all entities that appear in the ownership network and excluding isolates that never participate in a link. Each ownership relation—defined by a start and end date—is represented as a sequence of edge events. Specifically, we record a creation event at its start date and a deletion event at its termination.⁴ In addition, whenever the ownership share associated with an active relation changes over time, we record a link update event at the corresponding timestamp. This decomposition allows the model to track both the life cycle of ownership ties and the evolution of their associated attributes. To stabilize the input features, we apply a $\log(1+x)$ transformation to ownership shares and share capital, and we convert historical values from Italian Lira to Euros for temporal comparability.⁵

Firm-level financial statements are treated as node update events. Each firm generates one such event per year, represented as a self-loop (an edge connecting the node to itself) timestamped on December 31st of the reporting year. These self-loops inject updated firm-specific information—financial indicators and judicial status—into the model. Firms that are confiscated undergo a discrete feature change at the moment of seizure, when their status variable switches from “not confiscated” to the relevant judicial stage and their “confiscated share” variable updates from 0% to the seized percentage. Indi-

⁴When an end date for a specific pair corresponds to a start date for the same pair, we treat this as a link update rather than as a deletion followed by a creation.

⁵1 euro $\approx$ 1936.27 Italian Lira.

viduals (PFs), lacking such information, never generate node update events and evolve purely through their relational activity. This asymmetry is handled through a firm-level mask ensuring that node features are updated only for firms.

Finally, we merge link and node events into a single temporal stream, each event carrying a unified message vector of its features together with a binary flag distinguishing node updates from link events; events sharing a timestamp are ordered deterministically, with node deletions first, then link updates, then link creations. The sequence is split chronologically into training (1993 to end-2015), validation (end-2015 to early 2019) and test (2019 to 2024) sets. Updates dominate the stream, reflecting the dynamic nature of ownership relations, followed by creations and then deletions. Figures A.6 and A.7 in the Appendix report event counts by year and by type, and Figure A.3 their breakdown across splits.

Two model variants The description above corresponds to what we call the *full-data* variant of the model: every firm generates a node update event at the end of every year, carrying its current financial indicators together with a time-varying confiscation indicator that switches on from the firm’s confiscation year onward, so a firm’s judicial status is visible to the model only once it has become legally known. We also train a *first-year* variant, which differs in a single respect: each firm’s attributes are observed only once, in a node event stamped immediately before its first ownership event, and no judicial information ever enters the node features. After that initial observation the memory of a firm in the first-year variant is updated only by ownership events, so its embedding moves with network structure alone; in the full-data variant it also moves with the firm’s balance sheet. The distinction matters for the empirical application, whose outcomes are balance-sheet variables, and Section 5.1 explains which representation each exercise uses. Because the first-year event stream contains a single node event per firm rather than one per year, it is considerably shorter (about 100 thousand training events, against about 360 thousand for the full-data stream); the two variants are otherwise identical in architecture, losses, and training procedure.

Architecture and Training The Temporal Graph Network (TGN) (Rossi et al., 2020) models the ownership network as a continuous stream of timestamped events, each reflecting a structural or attribute-level change in the system. Every firm or individual is represented by a node equipped with a dynamic memory state $m_u(t)$ that summarizes its

past interactions and evolving structural role. The model processes events sequentially, updating these memory states to capture both short-term fluctuations and long-term relational dependencies.

The core pipeline operates as follows. When an ownership event occurs between two nodes u and v at time t , the model retrieves each node’s current memory and the memories of its recent temporal neighbors, combines these with edge attributes (e.g., ownership share) and encodings of the time elapsed since each prior interaction, and aggregates them through an attention mechanism⁶ into a context-aware embedding $z_u(t)$. Node memories are then updated via a gated recurrent unit, so that each node’s state reflects the cumulative history of its interactions. Firm-level financial statements enter as self-loop events—edges connecting a node to itself—which refresh the node’s memory without triggering the full message-passing pipeline, preserving the separation between structural changes (ownership edges) and internal attribute updates (balance-sheet data).

Three output heads operate on these embeddings, each specialized for a distinct behavioral dimension of ownership dynamics. A *link prediction* head estimates the probability that a relation $u \rightarrow v$ exists at time t , trained with binary cross-entropy against sampled negative edges; a *level regression* head predicts the current log ownership share of active links under an ℓ_1 loss; and a *delta regression* head predicts the change in log share between consecutive observations under a Huber loss. The heads share the same encoder but are activated selectively by event type: creations primarily drive the link head, updates trigger both regression heads, and removals are supervised through the delta head. Their contributions to the total loss are balanced dynamically, the weight on regression rising when a batch contains more updates and deletions. Two auxiliary losses regularize the embedding space: a Barlow Twins term (Zbontar et al., 2021) encouraging decorrelation across embedding dimensions, and a temporal consistency term (Trivedi et al., 2019) penalizing abrupt changes in node representations over time. Appendix B gives the full loss formulation, the dynamic weighting scheme, the attention formalism, the negative-sampling procedure, the scoring architectures and all hyperparameter choices.

⁶implemented using the `TransformerConv` operator from PyTorch Geometric (Fey and Lenssen, 2019), a graph-based variant of query-key-value attention (Vaswani et al., 2017).

Training proceeds for 20 epochs using the Adam optimizer with gradient clipping and adaptive learning-rate scheduling, node memories being carried forward across batches to preserve temporal continuity. Each variant is trained with three random seeds; Table A.4 in the Appendix reports test-set performance averaged across seeds, and the embeddings used in the analysis are those of the best-performing seed of each variant. The model is implemented in PyTorch (Paszke et al., 2019) and PyTorch Geometric (Fey and Lenssen, 2019) and comprises roughly 206k parameters. Figure B.16 in the Appendix illustrates the full event flow, message routing, and update cycle.

Embedding Extraction After training we extract yearly snapshots of node embeddings. Querying the model at arbitrary dates would risk leakage, so we instead reset its temporal state and stream the *entire* chronological event sequence through it in evaluation mode, updating each node’s memory exactly as it would have evolved during training but without gradient steps. We pause at a grid of January 1st timestamps and record all node embeddings at that instant. The embedding labelled year T therefore reflects only events occurring *before* January 1st of year T , preventing any leakage from future events. The extracted embeddings have 64 dimensions, and the procedure is applied separately to the full-data and to the first-year variant. These yearly snapshots form the basis for all subsequent analyses of infiltration risk and structural dynamics within the ownership network; all of them are conducted in the original 64-dimensional space.

5. Embedding Validation in the Corporate Network

Before using the embeddings to date transitions into a high-risk regime, we assess whether they recover economically meaningful and temporally ordered features of the corporate network. Four exercises follow. We first examine whether the geometry of the embedding space places firms confiscated only later closer to firms already confiscated, and summarize that geometry in a continuous, real-time measure of network-based risk, the Infiltration Proximity Index (Section 5.1); we then test whether that index forecasts subsequent confiscation out of sample, against the full set of firm-level observables (Section 5.2); we trace how latent positions move when a firm first connects to a counterparty confiscated only subsequently (Section 5.3); and we examine how pre-confiscation neighborhoods reorganize after one of their members is seized (Section 5.4). These are

descriptive exercises, evaluating geometric organization, predictive content, and sensitivity to local network change. Section 5.1 states which of the two representations of Section 4, full-data or first-year, each exercise uses.

5.1. Embedding Geometry and Network-Based Risk

We now ask how confiscation-related information is arranged inside the embedding space, and whether that arrangement can be summarized by a single continuous measure. The exercise is conducted in the original 64-dimensional space rather than on a low-dimensional projection.

Representation Section 4 describes two variants of the model, and the choice between them is dictated by the use to which each measure is put. The index constructed in this subsection is subsequently related to firm outcomes in Section 6, and those outcomes are themselves balance-sheet variables. Computing it on the full-data representation would make those regressions circular, since yearly financial statements enter that model as self-loop events refreshing each firm’s memory with current balance-sheet information. The index carried into the empirical application is therefore computed on the first-year variant, in which each firm’s characteristics are observed only once, in its first observed year, so that the temporal evolution of the representation is driven exclusively by ownership events. The validation exercises of this section are evaluated against a judicial event rather than a financial outcome, and no such concern arises: the confiscation direction, the forecasting exercise and the reconstructor analysis use the full-data representation, while the exposure exercise of Section 5.3 uses the first-year index, the one deployed in the application. Link-prediction accuracy is somewhat lower under the first-year restriction (Table A.4 in Appendix B) and the resulting space is geometrically distinct from the full-data one, but the two encode very similar relative positions: the rank correlation of firm-pair cosine distances across the two variants exceeds 0.91.


The confiscation direction We first construct a scalar summary of where a firm sits along the axis that separates confiscated from non-confiscated entities. Working with ℓ_2 -normalized embeddings, we compute the centroid of firms that are *already* confiscated at the time they are observed—that is, firm-year observations whose confiscation year is at or before the observation year—and the centroid of firms never confiscated anywhere in the sample. The *confiscation direction* is the unit vector pointing from the second

centroid to the first. Two restrictions keep this construction temporally honest. The direction is estimated only on firm-year observations up to 2012, so it does not use the geometry of the later period over which it is evaluated; and a firm contributes to the confiscated centroid only in the years after its own confiscation, so no firm enters that centroid before its own seizure. One qualification remains. The opposing centroid is formed from firms never confiscated anywhere in the sample, and that membership is a full-sample judgement: a firm confiscated after 2012 is excluded from it. The axis is therefore free of *timing* leakage—no seizure contributes to it before it occurs—but not entirely independent of full-sample confiscation labels. Each firm-year is then scored by the inner product between its normalized embedding and this fixed direction. Higher scores indicate closer alignment with the region of the space occupied by firms whose confiscation was already legally known at the time of measurement.

Figure 2 reports the distribution of this score in four snapshot years. Within each panel, firms present in that year are split into four groups by their confiscation status *relative to that year*: never confiscated in the sample; eventually confiscated, but more than three years later; confiscated within the following three years; and already confiscated by that year. Future confiscation serves as an ex-post label for the grouping. It enters the construction of the axis only through the composition of the never-confiscated centroid, as noted above, and plays no role in the timing of the score, which at any date reflects only seizures already recorded by then. The ordering is monotone and stable across snapshots. Already-confiscated firms score highest, as they must. Crucially, firms that will be confiscated within the next three years score well above never-confiscated firms and above firms whose confiscation lies further in the future, while never-confiscated firms sit markedly lower and are far more dispersed. Future-confiscated firms are therefore, on average, positioned closer to the contemporaneously known confiscated region several years before any judicial intervention. This is the central geometric result: the embedding space contains temporally ordered information associated with later confiscation, rather than merely separating firms once their status has become legally known. The distributions nonetheless overlap substantially, and a non-negligible mass of never-confiscated firms scores in the range typical of future-confiscated ones—firms that share relational characteristics with infiltrated entities without ever being subject to enforcement.

The Infiltration Proximity Index The confiscation direction reduces the geometry to a single axis defined by two centroids. For the forecasting exercise and for the empirical

Figure 2: Position along the Confiscation Direction, by Future Confiscation Status



Notes: One panel per calendar year; the unit is a firm present in that year. The score is the inner product between the firm's ℓ_2 -normalized embedding and a fixed unit vector running from the never-confiscated to the already-confiscated centroid. Both centroids are estimated once on observations up to 2012, counting a firm as confiscated only from its own confiscation year onward, so no firm enters the confiscated centroid before its own seizure. The opposing centroid is formed from firms never confiscated anywhere in the sample, so the direction is free of timing leakage but not of full-sample label information. Groups serve only as ex-post labels. Gray lines span the 5th–95th percentile, blue bars the interquartile range; black circles are medians, red diamonds means. Full-data representation.

application we require a measure that is local rather than global—one that reflects how densely a firm’s own latent neighborhood is populated by confiscated entities—and that uses, at each date, only the seizures already recorded by then. We therefore summarize each firm’s position with a kernel-weighted proximity score, which we label the *Infiltration Proximity Index* (IPI).

Let $\mathcal{N}_t$ denote the set of nodes with an embedding snapshot in year t , and let $\mathcal{C}_t \subseteq \mathcal{N}_t$ denote the set of entities judicially confiscated at or before year t . For any two nodes i and j observed in year t we compute the cosine distance between their embeddings,

$$d_{ij}(t) = 1 - \frac{z_i(t) \cdot z_j(t)}{\|z_i(t)\| \|z_j(t)\|},$$

and convert it into a similarity weight through a Gaussian kernel,

$$K_{ij}(t) = \exp\left(-\frac{d_{ij}(t)^2}{2\sigma_t^2}\right),$$

where the bandwidth σ_t is set adaptively, to the median across nodes of the distance to their 20-th nearest neighbor in year t . The IPI of firm i in year t is the share of that firm’s total kernel weight that is directed towards entities whose confiscation is already legally known:

$$\text{IPI}_i(t) = \frac{\sum_{j \in \mathcal{C}_t} K_{ij}(t)}{\sum_{j \in \mathcal{N}_t} K_{ij}(t)} \in [0, 1].$$

The index is high when the firms occupying a given firm’s latent neighborhood are disproportionately drawn from the already-confiscated set, and low when that neighborhood is dominated by entities that have not been confiscated by that date. It is recomputed from scratch in every year, with the reference set enlarged by the seizures recorded up to that year, so its variation over time reflects movement of the firm within the space, movement of the space around it, and the arrival of new confiscations in its neighborhood.

Three properties of this construction should be kept in mind when interpreting what follows. First, the IPI is a normalized similarity score, not a calibrated probability: a value of 0.4 does not mean a 40% chance of confiscation. Second, the reference set is time-respecting: a firm enters $\mathcal{C}_t$ only from its own confiscation year onward, so the index labelled year t uses no seizure recorded after t and is, in this sense, a real-time score. This is what allows it to serve as a forecasting input in Section 5.2. The sums

include the focal firm itself, whose self-similarity carries unit weight; for a firm not yet confiscated this term enters only the denominator, but from its confiscation year onward it enters the numerator too and pushes the index towards one irrespective of network position. Every use of the index in this paper therefore relies on pre-confiscation years only.⁷ Third, a high IPI is not proof that a firm is infiltrated: it says that the firm occupies a region of the ownership-embedding space densely populated by entities whose confiscation is already legally known.

Geometry and observed topology Embedding proximity might be a mechanical restatement of direct ownership adjacency, in which case nothing would be gained relative to reading the graph itself. Comparing each firm’s k nearest neighbors in embedding space with its one-hop neighbors in the observed graph shows otherwise: embedding neighborhoods recover graph neighbors far more often than chance—precision at $k = 5$ exceeds the random benchmark by roughly three orders of magnitude—yet the absolute overlap remains small, with recall below about 17% even at $k = 20$. Latent proximity is thus related to observed topology without reproducing one-hop adjacency: it aggregates relational history and higher-order structure that direct links alone do not encode. Appendix C reports the full precision and recall figures by year and by k .

5.2. Forecasting Future Confiscation

The second exercise asks whether the index forecasts confiscation *before* the judicial event, and how it compares with conventional firm-level observables. For a firm i we estimate the probability $p_i = f(\mathbf{X}_i)$ that it is confiscated h years later, with $h = 1, \dots, 4$ the number of years before confiscation at which the most recent information is measured.⁸ We compare two information sets: the full block of firm-level characteristics and financial indicators of Table A.1, at the reference year and five lags (roughly 560 features), and the IPI alone, at the reference year and four lags (five features), computed on the full-data representation with the time-respecting reference set. Following Ash

⁷The forecasting exercise measures the index $h \geq 1$ years before seizure; the exposure exercise of Section 5.3 is restricted to never-confiscated firms; reconnectors are positioned in the pre-confiscation window of Section 5.4; and the transition dates of Section 6 are assigned on pre-confiscation years.

⁸Non-confiscated firms are assigned a synthetic reference year by propensity-score matching on stable attributes, so that the two groups are aligned in calendar time. Appendix D describes the procedure, the chronological design, the classifiers and their tuning.

et al. (2025), each set is paired with logistic regression, random forest, LightGBM and XGBoost. The TGN is estimated on events through 2015, and each classifier is trained on reference years up to $2016 - h$ and tested on strictly later ones, so no embedding or seizure recorded after a reference year enters its forecast.

Two results stand out (Figure 3; Tables D.25–D.28). First, five numbers dominate several hundred: the IPI attains a higher out-of-sample ROC-AUC than the firm-level block at every horizon and for every classifier, 0.83–0.90 against 0.63–0.75, with a gap of about twenty points at $h = 3$ that does not narrow with the horizon and is fully available to logistic regression. Second, the two sets fail differently. Firm-level data paired with tree-based learners are comparatively precise but recover only a quarter to two-fifths of the confiscated firms; the IPI recovers more than three-quarters of them at every horizon and for every classifier, at a precision of about one-half to three-fifths, several times the base rate. Firm-level data produce few false positives and many false negatives; the index produces the reverse. Since confiscation is a delayed and incomplete record of infiltration, this is the desirable profile for a screening measure: a firm flagged by the index but never confiscated occupies a region populated by firms whose infiltration was later confirmed in court, and Figure 2 shows that many never-confiscated firms sit exactly there. Such firms are candidates for scrutiny, whereas the false negatives of the firm-level models are infiltrations that go undetected.


Two qualifications apply. The exercise demonstrates predictive content, not the onset of infiltration, and the index is computed on the full-data representation, so it is not free of balance-sheet information. And because the network is built outward from confiscation cases, the two classes share their environment, which makes the problem harder than in studies drawing firms from the general population (Cariello et al., 2024) and means that these figures are not an estimate of nationwide screening performance.

5.3. New Exposure to Future-Confiscated Firms

Forecasting and cross-sectional geometry show that the embeddings contain information associated with later confiscation. We next move from global geometry to a concrete change in the ownership network, and ask whether the representation responds coherently when a firm’s ownership environment changes in a specific way.

We restrict attention to firms never confiscated anywhere in the sample, so that the

Figure 3: Out-of-Sample ROC Curves: Firm-Level Data versus IPI, by Forecast Horizon



(a) Three-Year-Ahead Prediction ($h = 3$)

(b) Four-Year-Ahead Prediction ($h = 4$)

Notes: Out-of-sample ROC curves by classifier and information set; h is the number of years before confiscation at which the most recent information is measured. Green: firm-level characteristics and financial indicators at the reference year and five lags. Red: Infiltration Proximity Index at the reference year and four lags, computed on the full-data representation with the time-respecting reference set. Training uses reference years up to 2016 $- h$, testing strictly later ones. Shorter horizons and full metrics are in Appendix D.

focal firm's own judicial status never changes during the window. For each such firm i we identify the first year in which it has at least one direct (one-hop) counterparty j not confiscated at that date but confiscated at a strictly later one, set that year to event time zero, and trace firm i 's average IPI from five years before to five years after. Each focal firm contributes a single event, anchored at its *first* qualifying connection. The panel is unbalanced—a firm enters the average at a given event time only if observed in the corresponding calendar year—and each firm-year receives equal weight in the event-time mean. The index is the first-year IPI, the variant carried into the empirical application. The resulting sample comprises 1,592 focal firms.

Figure 4 reports the trajectory. Average risk is already rising before the connection, increases most sharply between $t - 1$ and $t + 1$, and remains elevated and gently increasing thereafter. The design is informative precisely because the counterparty is legally unmarked at event time zero: at the moment the tie forms, j carries no judicial status that the ownership network could have registered, and j does not yet belong to the reference set $\mathcal{C}_t$ of the index, so the contemporaneous movement can be attributed

neither to the model observing a known enforcement action nor to the mechanical arrival of a reference-set member in the focal firm’s neighborhood. What it indicates is that the representation reacts to the arrival of a relation with a firm whose *pre-confiscation* relational history already resembles confiscation-associated structures.


Three limits on the interpretation should be stated plainly. First, the increase is not the causal effect of forming the connection. Firms select into ownership relations, and the visible pre-trend—risk is climbing for several years before $t = 0$ —indicates that the event arrives in the middle of a broader change in the focal firm’s network trajectory rather than at its beginning. Second, the outcome is the raw IPI rather than a within-year percentile or a calendar-adjusted residual, and events are staggered across calendar years, so part of the common upward drift may reflect changes in the composition of the network and of the confiscated set over time rather than firm-level movement. Third, because the reference set is updated every year, the counterparty j enters $\mathcal{C}_t$ from its own confiscation year onward, which falls inside or beyond the post-event window; part of the continued rise after $t = 0$ therefore reflects the eventual arrival of j , and of other later-confiscated firms, in the reference set rather than further movement of the focal firm. The exercise is best read as a diagnostic of model behavior—the representation registers relationally risky ties as they arrive—rather than as evidence about the consequences of forming them.

5.4. Network Reconfiguration After a Neighbor’s Confiscation

The final exercise asks whether the embeddings capture relational continuity after a confiscation disrupts a local corporate neighborhood. Confiscation removes one node from the ownership network, but the entities that surrounded it remain. If the latent space encodes something about which firms belong to such environments, the firms that subsequently move into the vacated neighborhood should not be a random draw from the population of eligible outsiders.

Construction Consider a confiscated firm c seized in year $T_c \in [2000, 2020]$. Its *pre-confiscation neighborhood*, or “old web”, is the set of entities directly linked to c in the undirected yearly ownership graphs of the three preceding years, $T_c - 3$ through $T_c - 1$; it includes individuals as well as firms, since ownership links run through natural persons, and we retain only events whose old web contains at least two members. Over the three

Figure 4: IPI Around the First Connection to a Future-Confiscated Firm




Notes: Average IPI around a focal firm’s first direct connection to a *future-confiscated counterparty*—an entity not confiscated as of the year of the connection but confiscated later. Event time zero is the first year with at least one such counterparty at one hop; each firm contributes one event. The outcome is the raw IPI computed on the first-year representation with the time-respecting reference set; the band is ± 1 standard error.

years following seizure we identify every firm $r \neq c$ directly linked to at least one old-web member, restricting eligibility to legal entities not themselves confiscated at or before T_c . Each candidate is scored by the share of old-web members it reaches after T_c that it was *not* already linked to during the pre-window, so that pre-existing relationships do not count as new connectivity. Candidates with a positive score are *post-confiscation reconnectors*; the highest-scoring candidate at each event is that event’s *top reconstructor*. Figure 5 illustrates the construction.

We deliberately use “reconnectors” rather than “replacers”. What the data identify is the establishment of new ownership links with the surviving members of a neighborhood; we observe nothing about whether an entrant assumes the confiscated firm’s structural or economic function. For the same reason we describe the eligible pool as external firms not confiscated at the time of the event, rather than as “clean” firms: a firm that is never confiscated is not thereby known to be uninfiltated.

Result Direct reconnection by another future-confiscated firm is rare: in most events, the entrants forming new links with the old web are firms that are themselves never confiscated. The informative comparison is therefore not whether reconnectors are con-

Figure 5: Identifying Post-Confiscation Reconnectors



Notes: Illustration of the definitions, not an empirical example. The focal firm C is seized in T_c ; its old web is the set of entities directly linked to C over $T_c - 3$ to $T_c - 1$, firms (A, B) and individuals (P1, P2) alike. Eligible candidates—legal entities other than C , not confiscated at or before T_c —are scored by the share of old-web members they reach over $T_c + 1$ to $T_c + 3$ and were not already linked to before (dashed ties do not count). Positive scores define reconnectors, the highest score at each event the top reconnector.


fiscated, but where they sit in the embedding space relative to firms whose eventual status is known. Figure 6 places top reconnectors and the broader set of reconnectors alongside two reference groups drawn from the same calendar years: firms confiscated after the event, and firms never confiscated. Positions are measured by the within-year percentile rank of the confiscation-direction score of Section 5.1, computed among firms not already confiscated in that year, so that common calendar-time movements are differenced out.

The ordering is monotone. Future-confiscated firms rank highest, never-confiscated firms lowest, and reconnectors fall in between; within the reconnectors, those reconnecting to the largest share of the old web rank systematically above the broader set. The median top reconnector sits near the 81st within-year percentile, essentially the position of the median future-confiscated firm (the 82nd), against roughly the 59th for all reconnectors and the 46th for never-confiscated firms. Importantly, these positions are measured *before* the new links are formed: a candidate’s embedding is taken from the most recent year available in the pre-confiscation window $T_c - 3$ to $T_c - 1$, so the ranking reflects the firm’s latent position as of the period preceding both the seizure and its own subsequent reconnection. The result is therefore not an artifact of the model updating

in response to the very links that define the group.

This indicates that post-confiscation reconnection is not random with respect to the latent structure the TGN captures: the firms reconnecting most broadly to a disrupted neighborhood already resembled confiscation-associated firms more than a firm drawn at random from the same years did, even though most of them are never confiscated themselves. The comparison is with firms present in those years at large, not with the eligible candidates that could have reconnected but did not; it therefore establishes that reconnectors are drawn from a distinctive region of the space, not that they are selected from within the eligible pool. We read this as evidence of relational continuity—the environment surrounding a confiscated firm is repopulated by entities drawn from a structurally similar region of the ownership space. It does not establish that organized crime deliberately selects these entrants, that criminal control persists after enforcement, that a reconnector assumes the economic role of the confiscated firm, or that confiscation causally produces the observed links.

Figure 6: Embedding Position of Post-Confiscation Reconnectors



Notes: Within-year percentile rank of the confiscation-direction score (Section 5.1), computed among firms not already confiscated in the year of measurement; the unit is a firm attached to one confiscation event with $T_c \in [2000, 2020]$. The labels *Top Replacers* and *All Replacers* denote, respectively, top reconnectors and all reconnectors as defined in Figure 5. The reference groups are firms drawn at random from the same event years, not the eligible candidates that did not reconnect. Full-data representation.

6. Empirical Application: Infiltration & Firm Performance

6.1. Empirical Strategy


As an empirical application, we document which financial and operational patterns accompany the transition into a high-risk regime within our sample of firms, and how these patterns align with existing accounts of how criminal groups exploit firms. This allows us to characterize the high-risk firms we can observe.

Dating the transition Using the IPI empirically requires converting a firm’s continuous trajectory into a single date, and the distinctions this relies on are worth stating explicitly. Confiscation is observed but delayed, and it dates enforcement rather than involvement; infiltration is latent, and the moment at which it begins is never observed; the IPI is neither, being a continuous measure of structural proximity, in embedding space, to firms whose confiscation is already legally known. We date a firm’s transition as the *first large and persistent increase* in its first-year IPI. An increase in year t , $\Delta\text{IPI}_i(t) = \text{IPI}_i(t) - \text{IPI}_i(t-1)$, is *large* when it lies in the top decile of the positive year-on-year changes observed across firms in that calendar year; the cutoff is recomputed every year, which keeps the criterion comparable as the network and the confiscated set evolve. It is *persistent* when the index in both $t+1$ and $t+2$ remains above the firm’s own mean over the three preceding years, $t-3$ to $t-1$, which rules out spikes that revert. Three features keep the rule conservative. Only the *first* large increase is a candidate: if it cannot be dated, the rule never moves on to a later, more convenient one. Boundary cases excluded otherwise: a candidate with fewer than three observed years before it or fewer than two after it, with a year missing inside the window $t-3$ to $t+2$, or falling after 2020, yields no transition date. And for eventually confiscated firms the rule sees pre-confiscation years only, with the whole window required to precede the seizure, because from its confiscation year onward a firm’s index is mechanically inflated by its own entry into the reference set (Section 5.1); the yearly cutoffs are computed on not-yet-confiscated firm-years for the same reason.

The rule sorts firms into three groups. *Treated* firms are eventually confiscated firms with a dated transition: 574 of the 2,596 confiscated firms in the panel, the remainder

being mostly firms with too short a pre-confiscation series or whose first large increase lacks three prior years. *Never-treated* firms are firms that are never confiscated, never record a large increase, and whose index never enters the top quartile of its yearly distribution: 6,791 firms whose proximity to the confiscated region stays low throughout. All other firms are excluded from the baseline estimation sample; in particular, a confiscated firm is never used as a control, whatever the level of its index. What the procedure delivers is a retrospective date for a sharp and lasting change in embedding proximity. Figure 7 illustrates the classification for a single firm, confiscated in 2015 and dated as entering the high-risk regime in 2006: its index is flat near zero through 2005, rises in 2006 by more than that year’s cutoff, and stays above its pre-jump mean in 2007 and 2008. Later large increases are ignored because only the first is a candidate, and the years from confiscation onward play no role.

Figure 7: IPI Trajectory and Dated High-Risk Transition for a Single Confiscated Firm



Notes: Top panel: the firm’s first-year IPI. Blue: pre-confiscation years, the only ones the rule uses; grey: years from confiscation onward, ignored by the rule. Triangles mark year-on-year increases above that calendar year’s top-decile cutoff. The orange line is the firm’s mean IPI over $t-3$ to $t-1$, extended (dotted) over $t+1$ and $t+2$, the years in which the index must stay above it. The black dashed line marks the dated transition, the red dashed line the confiscation year. Bottom panel: number and composition of the firm’s first- and second-order connections in the yearly ownership graph.

We exploit variation across firms in the timing of this transition. From its transition year onward a firm is classified as having entered a high-risk regime. At each point in time, treated firms that have not yet experienced the transition (not-yet-treated) and the never-treated group serve as a comparison for firms that have already entered the high-risk regime. Differences across groups therefore reflect differences in the timing and incidence of inferred high-risk transitions, rather than differences in infiltration status per se.

For each outcome Y_{it} , we estimate an event-study specification of the form:

$$Y_{it} = \alpha_i + \gamma_t + \sum_{k \in \mathcal{K}, k \neq -1} \beta_k \mathbb{1}\{t - t_i^* = k\} + \varepsilon_{it}, \quad (1)$$

where α_i are firm fixed effects, γ_t are year fixed effects, t_i^* denotes the dated transition year of firm i , and $k = -1$ is the omitted reference period. The coefficients β_k trace the dynamic evolution of firm outcomes relative to the dated high-risk transition. In additional specifications, we also include province-by-year and sector-by-year fixed effects.

Estimation is implemented using the interaction-weighted estimator of [Sun and Abraham \(2021\)](#),⁹ which ensures that each event-time coefficient is constructed using comparisons between firms that have already transitioned and firms that have not yet transitioned (including firms that never transition) within the same relative period, and is robust to heterogeneity across transition cohorts.

By construction, transition to high-risk regimes involves firms targeted by organized criminal groups and ultimately confiscated, implying that these firms may already face distinct circumstances prior to the first IPI jump—such as liquidity stress, financing constraints, or other unobserved shocks that make them attractive vehicles for criminal activity. This is even more true as we have chosen to carefully select infiltrated firms whose risk-level is significantly higher than the rest of firms in sample. For these reasons, parallel trends assumptions are less likely to hold when comparing these units to the

⁹Standard TWFE in staggered settings can use already-transitioned firms as controls for later-transitioning ones, contaminating event-time coefficients even in a descriptive exercise: the [Sun and Abraham \(2021\)](#) estimator avoids these forbidden comparisons by restricting each event-time coefficient to contrasts between transitioned and not-yet-transitioned (or never-transitioned) firms.

never-confiscated firms in sample.

Given our restrictive definition of the transition date, our estimates should be interpreted as descriptive evidence on how the outcomes and balance sheets of these firms evolve around the moment at which their ownership position moves into the high-risk region identified by the embeddings-based risk signal.

6.2. Results

We begin by testing changes in scale and profitability around the first IPI jump. Figure 8 reports the dynamic responses of revenues, total costs, and pre-tax profits, with regression estimates reported in Appendix Table A.5. The event-time coefficients indicate a discrete increase in firm scale coinciding with the first IPI jump, with revenues and total costs rising in parallel. Pre-tax profits increase on impact but display wider confidence intervals at longer horizons, consistent with a transitory improvement rather than a persistent change in profit generation.

This pattern closely aligns with existing firm-level evidence on organized crime infiltration, which documents expansion without sustained profitability gains (e.g., [Mirenda et al., 2022](#); [Arellano-Bover et al., 2024](#)). Relative to this literature, our results show that the expansion in scale coincides sharply with the first observable increase in infiltration risk, suggesting that growth is concentrated around a discrete transition in firms' economic behavior rather than unfolding gradually.

Figure 8 also reports event-study estimates for receivables, short-term debt, and long-term debt, with aggregate estimates reported in Appendix Table A.6. Around the IPI jump, all three variables increase, indicating balance-sheet expansion accompanied by both larger claims and higher liabilities. The timing of these responses is consistent with an intensification of operating and financing activity after the inferred risk transition.

These findings are consistent with prior evidence documenting balance-sheet growth and increased indebtedness among infiltrated firms ([Mirenda et al., 2022](#); [Arellano-Bover et al., 2024](#)). While we cannot separately identify the underlying financing sources, our results show that receivables and liabilities expand jointly at the onset of heightened infiltration risk, pointing to a broad increase in financial flows rather than a shift in a

single credit channel.¹⁰

Appendix Figure A.8 presents the evolution of selected operating cost components around the IPI jump, with corresponding estimates in Appendix Table A.7. Operating costs increase in categories that plausibly allow greater contractual discretion, including services and the use of third-party assets, alongside higher labor costs.

This pattern corroborates descriptive evidence reported for confiscated firms in earlier work (Arellano-Bover et al., 2024). Relative to the existing literature, our results show that cost reallocation occurs sharply at the transition into the high-risk regime, rather than accumulating gradually over time. We remain agnostic about the underlying mechanisms driving these reallocations and interpret them as part of a broader reorganization of expenditures during post-jump expansion.

Appendix Figure A.8 also examines liquidity and balance-sheet composition around the IPI jump, with aggregate estimates reported in Appendix Table A.8. Following the jump, receivables relative to assets increase, while cash holdings relative to assets decline. The claims–liquidity congestion index rises persistently, indicating that balance-sheet expansion is associated with growing claims and reduced liquidity rather than cash accumulation.


This configuration extends prior evidence on the financial fragility of infiltrated firms by documenting how illiquidity emerges at the onset of heightened infiltration risk. Whereas benign credit expansion would typically ease liquidity constraints, the observed combination of balance-sheet growth and declining cash ratios suggests that firms become increasingly congested as financial intermediaries rather than accumulating liquid buffers.

6.3. Robustness Checks

Alternative Clustering In the baseline analysis, inference relies on clustering at the firm level. As a robustness check, we instead cluster at the level of static ownership network components, defined as weakly connected components (WCCs) of the ownership

¹⁰Due to data limitations, short- and long-term debt categories capture aggregate balance-sheet positions and do not allow us to distinguish between bank credit, trade credit, or other forms of external financing.

Figure 8: Scale expansion and balance-sheet growth around the IPI jump



(a) Scale expansion and profitability

(b) Debt and receivables growth

Notes: Sun & Abraham (2021) interaction-weighted event-study estimates with firm and year fixed effects. Coefficients are relative to $t = -1$ (omitted). Standard errors are clustered at the firm level. The sample includes all firms; at each event time k , firms that have already experienced the dated IPI transition are compared to firms that have not yet experienced it (including never-treated firms).

graph constructed using the set union of all observed links over time. This alternative accounts for potential correlation in outcomes across firms that are structurally connected through ownership ties, decreasing the number of clusters from roughly 15 thousand (i.e. all firms), to 850. All event studies and regression tables are re-estimated under this clustering scheme (Figures A.11–A.14 and Appendix Tables A.13–A.16). Results remain robust.

Leave-one-cohort-out Regression To assess whether the estimated patterns are driven by any single transition cohort, we implement a leave-one-cohort-out (LOCO) robustness check. We iteratively exclude one cohort at a time and re-estimate the event-study specification on the remaining sample, using the same outcomes and specification as in the baseline analysis. For each outcome, we summarize the distribution of the resulting estimates by reporting their mean and standard deviation across iterations, alongside the baseline 95% confidence interval obtained using the full set of cohorts (Table A.21). Across all outcome categories, the LOCO estimates are highly stable and tightly concentrated around the baseline results. The mean estimates lie well within the corresponding baseline confidence intervals, and the standard deviations are small relative to the magnitude of the estimates. This indicates that no single cohort disproportionately drives the results.

6.4. Discussion

This analysis allows us to paint a clear picture on the characteristics of the high-risk firms we are working with. Taken together, the evidence documents a coherent pattern in which high-risk firms during infiltration periods exhibit a discrete reorganization of scale, balance-sheet composition, and liquidity around the first IPI jump. This transition is characterized by a sharp expansion in scale, a simultaneous increase in receivables and liabilities, a reallocation of operating costs toward more discretionary categories, and a persistent decline in liquidity. These changes are not accompanied by sustained improvements in profitability or cash accumulation.

The volatility in revenues and costs and the high costs for services is in line with the idea that the most exposed infiltrated firms are exploited by criminal groups to support illicit and corruptive activities and/or for money laundering and fiscal frauds (Parbonetti, 2021). Interestingly, we find that this type of infiltrated firms, while being

those at highest risk of being identified and confiscated by public authorities, tend to remain in operation for several years rather than being active for very few years before being shut by the mafia groups themselves.

The observed configuration is consistent with an interpretation in which firms that are caught as infiltrated and confiscated operate as conduits for financial and real flows rather than as standard profit-maximizing or liquidity-accumulating entities. Operating volumes and balance-sheet positions expand in a way that resembles pass-through activity, while liquidity remains constrained. While we do not identify specific underlying channels—such as trade credit, procurement manipulation, or other forms of contractual intermediation—the observed patterns point to increased throughput rather than capital deepening or efficiency gains.

By focusing on a very specific type of infiltrated firms rather than the full universe of firms captured by the mafia, our approach adds a complementary perspective to the existing literature on organized crime infiltration, which documents that infiltrated firms experience growth in revenues and inputs without commensurate improvements in productivity, profitability, or financial stability (e.g., [Mirenda et al., 2022](#); [Arellano-Bover et al., 2024](#)). Rather than comparing infiltrated firms to an external group of non-infiltrated firms, we exploit variation in timing within a population of firms that are judicially confirmed to be infiltrated ex post. This design shows that the expansionary and congested balance-sheet dynamics emphasized in earlier work emerge sharply around the dated transition into the high-risk regime, rather than accumulating gradually over time. The alignment between our results and these prior findings is consistent with the IPI jump identifying an economically meaningful moment in these firms’ trajectories.

A final consideration concerns the magnitude of the estimated changes, which are substantially larger than those typically reported in the literature—roughly one order of magnitude in log-scale. This reflects key aspects of our empirical setting rather than implausibly large behavioral responses. At $t = -1$, median revenues among treated firms are approximately 24,000 euros (Table [A.5](#)), placing them firmly in the micro-firm category. A log-point increase of 3 from this baseline corresponds to a shift to roughly 480,000 euros—a substantial but not implausible transition for a firm moving from near-dormancy to active operations under criminal control. Importantly, these large expansions are not accompanied by sustained improvements in profitability or liquidity.

Instead, growth coincides with rising receivables, increasing liabilities, and declining cash holdings, reinforcing the interpretation of infiltrated firms as high-throughput conduits rather than as firms undergoing benign or efficiency-driven growth. The higher magnitude of estimated coefficients relative to previous studies is again consistent with the fact that our sample of treated firms include no false positives and, as per our embeddings computation, they are those belonging to the highest infiltration risk regime.

7. Concluding Remarks

This paper demonstrates how learned graph representations can be used to measure and study latent economic phenomena embedded in relational data. While embedding methods are now widely used in economics for text and other unstructured inputs, applications to dynamic network structures remain rare. We show that temporal graph embeddings provide a powerful and flexible framework for extracting economically meaningful information from evolving ownership networks.

Our argument proceeds in three stages. We first summarize the geometry of the learned ownership embeddings in a single real-time index, the Infiltration Proximity Index, and establish that it carries predictive content. In a chronological out-of-sample forecasting exercise, the index and four of its lags attain a higher area under the ROC curve than the full set of firm-level financial variables at every horizon up to four years before confiscation and for every classifier, and they do so by missing far fewer confiscated firms, at the cost of flagging more firms that enforcement never reaches—the profile a screening instrument should have when the underlying infiltration is observed only through delayed and incomplete enforcement. A natural concern regards the external validity of these results, given that the ownership network is constructed around judicially confirmed cases. This sample design implies that non-confiscated firms in our dataset are structurally and relationally proximate to infiltrated entities, making the discrimination problem substantially harder than comparing infiltrated firms to a random draw from the Italian corporate universe. Establishing how the framework would perform as a screening device on the universe of Italian ownership relations nonetheless requires access to those data; deployed at scale, methods of this kind could support authorities in prioritizing monitoring efforts and allocating investigative resources, without substituting judicial or investigative work.

We then validate the relational content of the representations directly. The geometry of the embedding space places firms that are confiscated only later closer to the region occupied by firms whose confiscation is already legally known at the time of measurement—a construction in which no seizure contributes to the confiscated centroid before it occurs. Latent positions move coherently when a firm first acquires a tie to a counterparty that is legally unmarked at the time but confiscated subsequently. And when the neighborhood of a confiscated firm is repopulated by previously external firms, the entrants reconnecting to the largest share of that neighborhood already occupied riskier positions in the space than the broader pool of entrants, measured before those links were formed. These exercises are descriptive: they document what the representations encode, not the effects of forming connections or of enforcement itself. The same index, computed on a variant of the model that observes firm attributes only in a firm’s first year, then serves as a timing device, dating when a firm’s ownership position first moves into a high-risk region of the latent space.

Finally, we use variation in the timing of IPI jumps—rather than exogenous treatment assignment—to trace which observable economic characteristics accompany these transitions. The evidence documents a sharp reorganization of firm activity: rapid expansion in scale, joint growth in receivables and liabilities, reallocation of operating costs toward more discretionary categories, and persistent liquidity deterioration. These patterns are consistent with firms operating as high-throughput conduits—intermediating financial and real flows—rather than accumulating capital, strengthening liquidity, or generating sustained profitability.

More broadly, this study illustrates how relational embeddings can be integrated with standard econometric frameworks to study latent exposure, structural risk, and dynamic transitions in complex economic systems. To the extent that criminal groups in other national contexts exploit corporate ownership through similar mechanisms—indirect chains, rapid turnover, temporary control—the framework proposed here could be adapted to detect infiltration risk beyond the Italian setting. More generally, the approach opens a data-driven pathway for analyzing environments in which influence and control propagate through evolving networks, extending the scope of embedding methods in applied economics.

References

- Arellano-Bover, J., De Simoni, M., Guiso, L., Macchiavello, R., Marchetti, D. J., and Prem, M. (2024). Mafias and firms.
- Ash, E., Chen, D. L., and Ornaghi, A. (2024). Gender attitudes in the judiciary: Evidence from us circuit courts. *American Economic Journal: Applied Economics*, 16(1):314–350.
- Ash, E., Galletta, S., and Giommoni, T. (2025). A machine learning approach to analyze and support anticorruption policy. *American Economic Journal: Economic Policy*, 17(2):162–193.
- Ash, E. and Hansen, S. (2023). Text algorithms in economics. *Annual Review of Economics*, 15:659–688.
- Ballester, C., Calvó-Armengol, A., and Zenou, Y. (2004). Who’s who in crime network. wanted the key player. Technical report, IUI Working Paper.
- Boeri, F., Di Cataldo, M., and Pietrostefani, E. (2024). Localized effects of confiscated and re-allocated real estate mafia assets. *Journal of Economic Geography*, 24(2):219–240.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.
- Brogaard, J. and Chen, B. C. (2024). Attention-based graph neural networks in firm cds prediction.
- Calamunci, F. and Drago, F. (2020). The economic impact of organized crime infiltration in the legal economy: Evidence from the judicial administration of organized crime firms. *Italian Economic Journal*, 6(2):275–297.
- Cariello, P., De Simoni, M., and Iezzi, S. (2024). A machine learning approach for the detection of firms infiltrated by organised crime in italy. *Quaderni dell’antiriciclaggio d’Italia, Unità di Informazione Finanziaria per l’Italia, Roma, Banca*.
- Castelluccio, M. and Rizzica, L. (2026). Mafia infiltrations in times of crisi. evidence from the covid-19 shock. *Journal of Law and Economics*, 69(3).
- Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. (2002). Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357.
- Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Pro-*

- ceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.
- Chircop, J., Fabrizi, M., Malaspina, P., and Parbonetti, A. (2023). Anti-mafia police actions, criminal firms, and peer firm tax avoidance. *Journal of Accounting Research*, 61(1):243–277.
- De Feo, G. and De Luca, G. D. (2017). Mafia in the ballot box. *American Economic Journal: Economic Policy*, 9(3):134–167.
- De Luca, R., Greco, R., Immordino, G., and Oliviero, T. (2025). Mafia infiltration and ownership dynamics in italian companies during covid-19. Technical report, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy.
- Dell, M. (2025). Deep learning for economists. *Journal of Economic Literature*, 63(1):5–58.
- Di Cataldo, M. and Mastrorocco, N. (2022). Organized crime, captured politicians, and the allocation of public resources. *The Journal of Law, Economics, and Organization*, 38(3):774–839.
- Europol (2024). Leveraging legitimacy: How the eu’s most threatening criminal networks abuse legal business structures. Technical report, Publications Office of the European Union. ISBN 978-92-95236-92-9, DOI: 10.2813/7731594, QL-01-24-014-EN-N.
- Fenzia, A. and Saggio, R. (2024). Organized crime and economic growth: Evidence from municipalities infiltrated by the mafia. *American Economic Review*, 114(7):2171–2200.
- Fey, M. and Lenssen, J. E. (2019). Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Gentzkow, M., Kelly, B., and Taddy, M. (2019a). Text as data. *Journal of Economic Literature*, 57(3):535–574.
- Gentzkow, M., Shapiro, J. M., and Taddy, M. (2019b). Measuring group differences in high-dimensional choices: method and application to congressional speech. *Econometrica*, 87(4):1307–1340.
- Gorodnichenko, Y., Pham, T., and Talavera, O. (2023). The voice of monetary policy. *American Economic Review*, 113(2):548–584.
- Hansen, S., McMahon, M., and Prat, A. (2018). Transparency and deliberation within the fomc: A computational linguistics approach. *The Quarterly Journal of Economics*, 133(2):801–870.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y.


- (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.
- Le Moglie, M. and Sorrenti, G. (2022). Revealing “mafia inc.”? financial crisis, organized crime, and the birth of new enterprises. *Review of Economics and Statistics*, 104(1):142–156.
- Lindquist, M. J. and Zenou, Y. (2019). Crime and networks: Ten policy lessons. *Oxford Review of Economic Policy*, 35(4):746–771.
- Lu, C., Durante, D., and Friel, N. (2025). Zero-inflated stochastic block modelling of efficiency-security trade-offs in weighted criminal networks. *Journal of the Royal Statistical Society Series A: Statistics in Society*, page qnaf029.
- Magnolfi, L., McClure, J., and Sorensen, A. (2025). Triplet embeddings for demand estimation. *American Economic Journal: Microeconomics*, 17(1):282–307.
- Mirenda, L., Mocetti, S., and Rizzica, L. (2022). The economic effects of mafia: Firm level evidence. *American Economic Review*, 112(8):2748–2773.
- Parbonetti, A. (2021). La presenza delle mafie nell’economia: profili e modelli operativi. In *La presenza delle mafie nell’economia: profili e modelli operativi: Parbonetti, Antonio*. Padova: Padova University Press.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32.
- Pinotti, P. (2015). The economic costs of organised crime: Evidence from southern Italy. *The Economic Journal*, 125(586):F203–F232.
- Pinotti, P., Daniele, G., De Simoni, M., Marchetti, D., and Marcolongo, G. (2025). *A Loan You Can’t Refuse: Credit Rationing and Organized Crime Infiltration of Distressed Firms*. RFBerlin, Rockwool Foundation Berlin.
- Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., and Bronstein, M. (2020). Temporal graph networks for deep learning on dynamic graphs. *arXiv preprint arXiv:2006.10637*.
- Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199.
- Trivedi, R., Farajtabar, M., Biswal, P., and Zha, H. (2019). Dyrep: Learning representa-

- tions over dynamic graphs. In *International Conference on Learning Representations (ICLR)*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.
- Villarraga, D. F. and Daziano, R. A. (2025). Designing graph convolutional neural networks for discrete choice with network effects. *arXiv preprint arXiv:2503.09786*.
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Dévlin, S. (2021). Barlow twins: Self-supervised learning via redundancy reduction. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pages 12310–12320. PMLR.
- Zheng, Y., Xu, Z., and Xiao, A. (2023). Deep learning in economics: a systematic and critical review. *Artificial Intelligence Review*, 56(9):9497–9539.

A. Appendix


A.1. Figures

Figure A.1: Distribution of Confiscations over time




Notes: The figure shows the distribution of confiscation dates over time as reported by ANBSC. The confiscation date indicates the beginning of legal proceedings against the firm.

Figure A.2: Geographical Distribution of Confiscations




Notes: The figure shows the municipality-level geographic distribution of confiscations over Italy as reported by ANBSC. The size of the circle indicates the number of firms.

Figure A.3: Sector Distribution of Confiscated and Non-Confiscated Firms




Notes: The figure shows the sectorial distribution of confiscated and non-confiscated firms, following the NACE classification.

Figure A.4: Static Graph




(a) All Nodes



(b) Main Component

Notes: Panel (a) shows the static graph, including all nodes and edges over time. Panel (b) highlights nodes belonging to the main component. No node coloring for the static graph.


Figure A.5: Network Snapshots over Time




(a) Network Snapshot: 1998



(b) Network Snapshot: 2001




(c) Network Snapshot: 2010



(d) Network Snapshot: 2016


Notes: Different snapshots of the dynamic graph, showcasing visualization features. In panel A.5c the right banner displays node attributes and the left banner edge attributes. Panel A.5d displays the main connected component with the three hub nodes

Figure A.6: Network Events Over Time




Notes: The figure illustrates the number of events per year, considering both edge events, i.e. *interaction events* in blue and node events, i.e. *self-loop events* in orange.

Figure A.7: Edge Events over Time



Notes: The figure shows the breakdown of edge events by type, distinguishing between link creation (in green), link deletion (in red).

Figure A.8: Cost reallocation and liquidity dynamics around the IPI jump




(a) Cost reallocation

(b) Liquidity and balance-sheet congestion


Notes: Sun & Abraham (2021) interaction-weighted event-study estimates with firm and year fixed effects. Coefficients are relative to $t = -1$ (omitted). Standard errors are clustered at the firm level. The sample includes only eventually-confiscated firms; at each event time k , firms that have already experienced the dated IPI transition are compared to firms that have not yet experienced it (including never-treated firms).

Figure A.9: Scale expansion and balance-sheet growth around the IPI jump - confiscated only



Notes: Sun & Abraham (2021) interaction-weighted event-study estimates with firm and year fixed effects. Coefficients are relative to $t = -1$ (omitted). Standard errors are clustered at the firm level. The sample includes all firms; at each event time k , firms that have already experienced the dated IPI transition are compared to firms that have not yet experienced it (including never-treated firms).

Figure A.10: Cost reallocation and liquidity dynamics around the IPI jump - confiscated only




(a) Cost reallocation


(b) Liquidity and balance-sheet congestion

Notes: Sun & Abraham (2021) interaction-weighted event-study estimates with firm and year fixed effects. Coefficients are relative to $t = -1$ (omitted). Standard errors are clustered at the firm level. The sample includes only eventually-confiscated firms; at each event time k , firms that have already experienced the dated IPI transition are compared to firms that have not yet experienced it (including never-treated firms).


Figure A.11: Inference Robustness: Scale expansion and profitability around the IPI jump



(a) Sales and operating revenues




(b) Total production costs




(c) Pre-tax profit

Notes: Sun & Abraham (2021) interaction-weighted event-study estimates with firm and year fixed effects. Coefficients are relative to $t = -1$ (omitted). Standard errors are clustered at the static ownership graph's weakest connected component level.


Figure A.12: Inference Robustness: Receivables and debt dynamics around the IPI jump



(a) Total receivables




(b) Short-term debt




(c) Long-term debt

Notes: Sun & Abraham (2021) interaction-weighted event-study estimates with firm and year fixed effects. Coefficients are relative to $t = -1$ (omitted). Standard errors are clustered at the static ownership graph's weakest connected component level.


Figure A.13: Inference Robustness: Cost components around the IPI jump



(a) Services




(b) Use of third-party assets




(c) Labor costs

Notes: Sun & Abraham (2021) interaction-weighted event-study estimates with firm and year fixed effects. Coefficients are relative to $t = -1$ (omitted). Standard errors are clustered at the static ownership graph's weakest connected component level.


Figure A.14: Inference Robustness: Liquidity and balance-sheet congestion around the IPI jump



(a) Receivables-to-assets ratio



(b) Cash-to-assets ratio



(c) Claims-liquidity congestion index

Notes: Sun & Abraham (2021) interaction-weighted event-study estimates with firm and year fixed effects. Coefficients are relative to $t = -1$ (omitted). Standard errors are clustered at the static ownership graph's weakest connected component level.

A.2. Tables

Table A.1: List of Firm-level Variables from *Infocamere* by Category

Firm Characteristics	Financial Indicators	Balance Sheet: Assets	Balance Sheet: Liabilities & Equity	Income Statement
size	i_valpro	aa_crediti_verso_soci	ba_totale_patrimonio_netto	ea_totale_valore_della_produzione
status	i_valagg	ab_totale_immobilizzazioni	ba1_capitale_sociale	ea01_ricavi_vendite_e_prestazioni
comune	i_marope	ab1_totale_immobilizzazioni_immateriali	ba2_riserva_da_sovraprezzo	ea02_variazioni_rimanenze_prodotti
provincia	i_risuope	ab2_totale_immobilizzazioni_materiali	ba3_riserva_di_rivalutazione	ea03_variazione_lavori_in_corso_sul_ordinazione
regione	i_risugca	ab3_totale_immobilizzazioni_finanziarie	ba4_riserva_legale	ea04_incrementi_di_immobilizz_per_lav_interni
n_addetti_tot	i_risuaim	ac_attivo_circolante	ba5_riserva_statutaria	ea05_altri_ricavi
status_attivita	i_risumet	ac1_totale_rimanenze	ba6_riserva_azioni_proprie	eb_totale_costi_della_produzione
c_s_attivita	i_immob	ac2_totale_crediti	ba7_altre_riserve	eb06_acquisti_di_materie
c_attivita	i_rimane	ac3_totale_attivita_finanziarie	ba8_utile_perdita_a_nuovo	eb07_servizi
c_ateco_benc	i_crediti	ac4_disponibilita_liquide	ba9_utile_perdita_di_esercizio	eb08_godimento_di_beni_di_terzi
gruppo_impresa	i_displiq	ad_ratei_e_risconti	bax_capitale_proprio	eb09_costi_per_il_personale
ateco_letter	i_totatt	at_totale_attivo	bb_totale_fondi_rischi	eb10_ammortamenti_e_svalutazioni
industry	i_patrnet		bb1_fondo_di_quiescenza	eb11_variazione_rimanenze_materie_prime
	i_pascous		bb2_fondo_imposte	eb12_accantonamenti_per_rischi
	i_pascorr		bb3_altri_fondi	eb13_altri_accantonamenti
	i_totpass		bc_trattamento_di_fine_rapporto	eb14_oneri_diversi_di_gestione
	i_roe		bd_totale_debiti	ec_totale_proventi_e_oneri_finanziari
	i_roi		bd1_obbligazioni	ec15_proventi_da_partecipazioni
	i_indfin		bd10_debiti_verso_imprese_collegate	ec16_altri_proventi_finanziari
	i_copimm		bd11_debiti_verso_controllanti	ec17_interessi_e_altri_oneri_finanziari
	i_liqimm		bd2_obbligazioni_convertibili	ed18_rivalutazioni
			bd4_totale_debiti_vs_banche	ed19_svalutazioni
			bd5_totale_debiti_vs_altri_finanziatori	ee_totale_proventi_e_oneri_straordinari
			bd6_acconti_anticipi_da_clienti	ef00_risultato_prima_delle_imposte
			bd7_totale_debiti_vs_fornitori	ef22_imposte_sul_reddito
			bd9_debiti_verso_imprese_controllate	ef23_utile_perdita_di_esercizio
			bdx_totale_debiti_commerciali	
			bdy_totale_debiti_residui	
			bdz_a_totale_debiti_entro_esercizio	
			bdz_b_totale_debiti_oltre_esercizio	
			be_ratei_e_risconti	
			bt_totale_passivo	

Notes: The table reports the firm-level features used for the predictive modeling, categorized into five groups: Firm Characteristics, Financial Indicators, Balance Sheet Assets, Balance Sheet Liabilities & Equity, and Income Statement. Each variable is represented by its code as per the *Infocamere* database.

Table A.2: Edge Event Breakdown and Link Duration Statistics

Metric	Value
Total edge events	264,284
Link creations	50,685
Link updates	163,683
Link deletions	49,916
Creation/Deletion ratio	1.02
Updates/Creations ratio	3.23
Links with complete lifecycle	49,910
Average link duration (days)	1859.07
Std deviation (days)	1808.20
Average link duration (years)	5.09
Std deviation (years)	4.95

Notes: Link creations represent new connections established in the network. Link updates indicate modifications to existing connections (e.g., changes in relationship type or strength). Link deletions occur when connections are dissolved or removed from the network. Link duration is calculated only for links with both creation and deletion events, representing the complete lifecycle from establishment to dissolution. Duration statistics are computed in both days and years, where 1 year = 365.25 days.

Table A.3: Event Type Distribution Across Data Splits

Event Type	Train	Validation	Test	Total
Node events	264,835	62,140	68,305	395,280
Link creations	43,479	5,036	2,170	50,685
Link updates	122,805	22,981	17,897	163,683
Link deletions	31,329	8,585	10,002	49,916
Total events	462,448	98,742	98,374	659,564

Notes: Node events represent updates to firm-level attributes (e.g., financial data). Link creations are new connections established between firms. Link updates modify existing connections. Link deletions remove connections from the network. The data is split temporally into train, validation, and test sets for model evaluation.

Table A.4: Performance Comparison of TGN Model Variants

Metric	Full-data TGN	First-year TGN
<i>Link Prediction</i>		
AP	0.832 ± 0.013	0.743 ± 0.023
AUC	0.981 ± 0.002	0.964 ± 0.006
<i>Level Regression (creations + updates)</i>		
MAE	0.712 ± 0.032	0.770 ± 0.049
RMSE	1.022 ± 0.038	1.093 ± 0.068
<i>Delta Regression (updates + deletions)</i>		
MAE_Δ	0.724 ± 0.017	0.721 ± 0.015
RMSE_Δ	1.198 ± 0.021	1.181 ± 0.004
Training events	361,663	99,891
Test events	98,340	27,875

Note: Test-set metrics averaged over 3 random seeds (42, 123, 456); $\pm$ denotes the standard deviation across seeds. Both variants share the same architecture (64-dimensional memory, MLP scoring head, about 206k parameters) and are trained for 20 epochs. Full-data TGN: yearly firm-level attributes and a time-varying confiscation indicator enter as node-update events stamped at year end. First-year TGN: firm-level attributes are injected once, just before the firm’s first ownership event, and no judicial information enters; the event stream is correspondingly shorter. AP: Average Precision, AUC: Area Under ROC Curve. MAE/RMSE: regression metrics for absolute log stake levels (creations+updates). MAE_Δ/RMSE_Δ: regression metrics for log stake changes (updates+deletions).

Table A.5: Scale Expansion and Profitability

	(1)	(2)	(3)	(4)
<i>Panel A: Revenues (IHS)</i>				
INF_{it}	3.606*** (0.563)	3.588*** (0.582)	3.575*** (0.568)	3.530*** (0.584)
Mean	632,060.3	632,060.3	632,060.3	632,060.3
Median	11,916	11,916	11,916	11,916
N	121,813	121,813	121,813	121,813
<i>Panel B: Total Costs (IHS)</i>				
INF_{it}	1.857*** (0.220)	1.771*** (0.223)	1.828*** (0.224)	1.721*** (0.226)
Mean	672,946	672,946	672,946	672,946
Median	45,732	45,732	45,732	45,732
N	121,813	121,813	121,813	121,813
<i>Panel C: Pre-Tax Profit (IHS)</i>				
INF_{it}	2.758** (1.133)	2.880** (1.159)	2.598** (1.136)	2.674** (1.162)
Mean	-36,752.46	-36,752.46	-36,752.46	-36,752.46
Median	-6	-6	-6	-6
N	121,813	121,813	121,813	121,813
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Debt and Receivables Growth

	(1)	(2)	(3)	(4)
<i>Panel A: Receivables (IHS)</i>				
INF_{it}	1.948*** (0.302)	1.870*** (0.304)	1.899*** (0.304)	1.823*** (0.305)
Mean	322,097.6	322,097.6	322,097.6	322,097.6
Median	22,315	22,315	22,315	22,315
N	121,251	121,251	121,251	121,251
<i>Panel B: Short-Term Debt (IHS)</i>				
INF_{it}	1.565*** (0.299)	1.493*** (0.306)	1.524*** (0.300)	1.463*** (0.306)
Mean	526,676.7	526,676.7	526,676.7	526,676.7
Median	65,568	65,568	65,568	65,568
N	121,251	121,251	121,251	121,251
<i>Panel C: Long-Term Debt (IHS)</i>				
INF_{it}	1.719*** (0.502)	1.797*** (0.509)	1.683*** (0.504)	1.769*** (0.512)
Mean	176,271	176,271	176,271	176,271
Median	0	0	0	0
N	121,251	121,251	121,251	121,251
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Cost Reallocation

	(1)	(2)	(3)	(4)
<i>Panel A: Services (IHS)</i>				
INF_{it}	2.041*** (0.334)	1.926*** (0.340)	2.037*** (0.336)	1.907*** (0.342)
Mean	239,882	239,882	239,882	239,882
Median	5,891	5,891	5,891	5,891
N	121,813	121,813	121,813	121,813
<i>Panel B: Third-Party Assets (IHS)</i>				
INF_{it}	1.756*** (0.419)	1.641*** (0.432)	1.715*** (0.421)	1.581*** (0.432)
Mean	28,769.84	28,769.84	28,769.84	28,769.84
Median	0	0	0	0
N	121,813	121,813	121,813	121,813
<i>Panel C: Labor Costs (IHS)</i>				
INF_{it}	2.805*** (0.440)	2.817*** (0.439)	2.808*** (0.439)	2.789*** (0.437)
Mean	61,024.41	61,024.41	61,024.41	61,024.41
Median	0	0	0	0
N	121,813	121,813	121,813	121,813
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.8: Claims Accumulation and Liquidity Drain

	(1)	(2)	(3)	(4)
<i>Panel A: Receivables/Assets</i>				
INF_{it}	0.103*** (0.025)	0.088*** (0.025)	0.093*** (0.025)	0.080*** (0.025)
Mean	0.28	0.28	0.28	0.28
Median	0.18	0.18	0.18	0.18
N	121,170	121,170	121,170	121,170
<i>Panel B: Cash/Assets</i>				
INF_{it}	-0.044* (0.024)	-0.038 (0.024)	-0.042* (0.024)	-0.036 (0.024)
Mean	0.14	0.14	0.14	0.14
Median	0.06	0.06	0.06	0.06
N	121,732	121,732	121,732	121,732
<i>Panel C: Claims-Liquidity Congestion</i>				
INF_{it}	0.148*** (0.042)	0.127*** (0.041)	0.135*** (0.042)	0.117*** (0.041)
Mean	0.14	0.14	0.14	0.14
Median	0.1	0.1	0.1	0.1
N	121,170	121,170	121,170	121,170
<i>Panel D: Financial Share of Liabilities</i>				
INF_{it}	-0.113** (0.055)	-0.120* (0.066)	-0.224** (0.088)	-0.191* (0.103)
Mean	-0.16	-0.16	-0.16	-0.16
Median	-0.14	-0.14	-0.14	-0.14
N	121,230	121,230	121,230	121,230
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. *p<0.1, **p<0.05, ***p<0.01.

Table A.9: Scale Expansion and Profitability (Confiscated Only)

	(1)	(2)	(3)	(4)
<i>Panel A: Revenues (IHS)</i>				
INF_{it}	3.270*** (0.586)	2.783 (39.159)	3.583*** (0.599)	3.423 (96.864)
Mean	632,060.3	632,060.3	632,060.3	632,060.3
Median	11,916	11,916	11,916	11,916
N	4,241	4,241	4,241	4,241
<i>Panel B: Total Costs (IHS)</i>				
INF_{it}	1.674*** (0.235)	1.618 (16.204)	1.777*** (0.247)	1.743 (47.594)
Mean	672,946	672,946	672,946	672,946
Median	45,732	45,732	45,732	45,732
N	4,241	4,241	4,241	4,241
<i>Panel C: Pre-Tax Profit (IHS)</i>				
INF_{it}	2.926** (1.178)	1.751 (145.910)	2.806** (1.251)	2.309 (217.058)
Mean	-36,752.46	-36,752.46	-36,752.46	-36,752.46
Median	-6	-6	-6	-6
N	4,241	4,241	4,241	4,241
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.10: Debt and Receivables Growth (Confiscated Only)

	(1)	(2)	(3)	(4)
<i>Panel A: Receivables (IHS)</i>				
INF_{it}	1.650*** (0.314)	1.385 (6.315)	1.506*** (0.365)	1.254 (20.088)
Mean	322,097.6	322,097.6	322,097.6	322,097.6
Median	22,315	22,315	22,315	22,315
N	4,209	4,209	4,209	4,209
<i>Panel B: Short-Term Debt (IHS)</i>				
INF_{it}	1.245*** (0.312)	1.174 (8.502)	1.299*** (0.328)	1.251 (27.975)
Mean	526,676.7	526,676.7	526,676.7	526,676.7
Median	65,568	65,568	65,568	65,568
N	4,209	4,209	4,209	4,209
<i>Panel C: Long-Term Debt (IHS)</i>				
INF_{it}	1.242** (0.537)	1.426 (18.223)	0.862 (0.586)	1.230 (54.169)
Mean	176,271	176,271	176,271	176,271
Median	0	0	0	0
N	4,209	4,209	4,209	4,209
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.11: Cost Reallocation (Confiscated Only)

	(1)	(2)	(3)	(4)
<i>Panel A: Services (IHS)</i>				
INF_{it}	1.852*** (0.347)	1.508 (27.523)	1.859*** (0.359)	1.740 (38.200)
Mean	239,882	239,882	239,882	239,882
Median	5,891	5,891	5,891	5,891
N	4,241	4,241	4,241	4,241
<i>Panel B: Third-Party Assets (IHS)</i>				
INF_{it}	1.471*** (0.438)	1.226 (35.089)	1.479*** (0.464)	1.326 (59.711)
Mean	28,769.84	28,769.84	28,769.84	28,769.84
Median	0	0	0	0
N	4,241	4,241	4,241	4,241
<i>Panel C: Labor Costs (IHS)</i>				
INF_{it}	2.451*** (0.461)	2.315 (29.726)	2.738*** (0.467)	2.572 (53.399)
Mean	61,024.41	61,024.41	61,024.41	61,024.41
Median	0	0	0	0
N	4,241	4,241	4,241	4,241
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.12: Claims Accumulation and Liquidity Drain (Confiscated Only)

	(1)	(2)	(3)	(4)
<i>Panel A: Receivables/Assets</i>				
INF_{it}	0.092*** (0.026)	0.079 (1.216)	0.080*** (0.029)	0.062 (3.369)
Mean	0.28	0.28	0.28	0.28
Median	0.18	0.18	0.18	0.18
N	4,208	4,208	4,208	4,208
<i>Panel B: Cash/Assets</i>				
INF_{it}	-0.032 (0.024)	-0.017 (1.659)	-0.039 (0.024)	-0.031 (3.703)
Mean	0.14	0.14	0.14	0.14
Median	0.06	0.06	0.06	0.06
N	4,240	4,240	4,240	4,240
<i>Panel C: Claims-Liquidity Congestion</i>				
INF_{it}	0.124*** (0.043)	0.098 (2.026)	0.119*** (0.045)	0.087 (6.477)
Mean	0.14	0.14	0.14	0.14
Median	0.1	0.1	0.1	0.1
N	4,208	4,208	4,208	4,208
<i>Panel D: Financial Share of Liabilities</i>				
INF_{it}	-4.289 (3.419)	-1.403 (96.571)	-7.947 (6.870)	-7.218 (347.623)
Mean	-0.17	-0.17	-0.17	-0.17
Median	-0.24	-0.24	-0.24	-0.24
N	4,208	4,208	4,208	4,208
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.13: Scale Expansion and Profitability (Inference Robustness)

	(1)	(2)	(3)	(4)
<i>Panel A: Revenues (Log)</i>				
INF_{it}	3.079*** (0.310)	3.108*** (0.275)	3.041*** (0.303)	3.064*** (0.272)
Mean	531,879.5	531,879.5	531,879.5	531,879.5
Median	23,737	23,737	23,737	23,737
N	139,600	139,600	139,600	139,600
<i>Panel B: Total Costs (Log)</i>				
INF_{it}	1.756*** (0.164)	1.764*** (0.158)	1.739*** (0.161)	1.740*** (0.155)
Mean	583,782.9	583,782.9	583,782.9	583,782.9
Median	49,453	49,453	49,453	49,453
N	139,600	139,600	139,600	139,600
<i>Panel C: Pre-Tax Profit (IHS)</i>				
INF_{it}	1.931*** (0.589)	1.942*** (0.613)	1.871*** (0.592)	1.845*** (0.616)
Mean	17,813.13	17,813.13	17,813.13	17,813.13
Median	-0.5	-0.5	-0.5	-0.5
N	139,600	139,600	139,600	139,600
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors are clustered at the static ownership graph's weakest connected component level. Mean and Median computed at $t = -1$ for treated firms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.14: Debt and Receivables Growth (Inference Robustness)

	(1)	(2)	(3)	(4)
<i>Panel A: Receivables (Log)</i>				
INF_{it}	1.904*** (0.140)	1.841*** (0.141)	1.864*** (0.140)	1.809*** (0.140)
Mean	233,674.8	233,674.8	233,674.8	233,674.8
Median	23,782.5	23,782.5	23,782.5	23,782.5
N	139,035	139,035	139,035	139,035
<i>Panel B: Short-Term Debt (Log)</i>				
INF_{it}	1.584*** (0.179)	1.515*** (0.178)	1.541*** (0.179)	1.481*** (0.179)
Mean	464,711.6	464,711.6	464,711.6	464,711.6
Median	87,457	87,457	87,457	87,457
N	139,035	139,035	139,035	139,035
<i>Panel C: Long-Term Debt (Log)</i>				
INF_{it}	1.960*** (0.359)	2.016*** (0.391)	1.897*** (0.372)	1.963*** (0.404)
Mean	190,997	190,997	190,997	190,997
Median	0	0	0	0
N	139,035	139,035	139,035	139,035
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors are clustered at the static ownership graph's weakest connected component level. Mean and Median computed at $t = -1$ for treated firms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.15: Cost Reallocation (Inference Robustness)

	(1)	(2)	(3)	(4)
<i>Panel A: Services (Log)</i>				
INF_{it}	1.960*** (0.227)	1.957*** (0.225)	1.933*** (0.223)	1.924*** (0.222)
Mean	101,167.5	101,167.5	101,167.5	101,167.5
Median	6,797	6,797	6,797	6,797
N	139,600	139,600	139,600	139,600
<i>Panel B: Third-Party Assets (Log)</i>				
INF_{it}	2.246*** (0.188)	2.216*** (0.187)	2.207*** (0.187)	2.181*** (0.188)
Mean	19,359.42	19,359.42	19,359.42	19,359.42
Median	0	0	0	0
N	139,600	139,600	139,600	139,600
<i>Panel C: Labor Costs (Log)</i>				
INF_{it}	2.810*** (0.210)	2.817*** (0.206)	2.739*** (0.203)	2.742*** (0.203)
Mean	53,901.87	53,901.87	53,901.87	53,901.87
Median	0	0	0	0
N	139,600	139,600	139,600	139,600
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors are clustered at the static ownership graph's weakest connected component level. Mean and Median computed at $t = -1$ for treated firms. *p<0.1, **p<0.05, ***p<0.01.

Table A.16: Claims Accumulation and Liquidity Drain (Inference Robustness)

	(1)	(2)	(3)	(4)
<i>Panel A: Receivables/Assets</i>				
INF_{it}	0.103*** (0.012)	0.092*** (0.012)	0.098*** (0.012)	0.088*** (0.012)
Mean	0.26	0.26	0.26	0.26
Median	0.14	0.14	0.14	0.14
N	138,937	138,937	138,937	138,937
<i>Panel B: Cash/Assets</i>				
INF_{it}	-0.046*** (0.009)	-0.044*** (0.009)	-0.045*** (0.009)	-0.042*** (0.009)
Mean	0.16	0.16	0.16	0.16
Median	0.06	0.06	0.06	0.06
N	139,502	139,502	139,502	139,502
<i>Panel C: Claims-Liquidity Congestion</i>				
INF_{it}	0.149*** (0.017)	0.135*** (0.017)	0.142*** (0.017)	0.130*** (0.018)
Mean	0.1	0.1	0.1	0.1
Median	0.05	0.05	0.05	0.05
N	138,937	138,937	138,937	138,937
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. Standard errors are clustered at the static ownership graph's weakest connected component level. Mean and Median computed at $t = -1$ for treated firms. *p<0.1, **p<0.05, ***p<0.01.

Table A.17: Scale Expansion and Profitability (Placebo Test)

	(1)	(2)	(3)	(4)
<i>Panel A: Revenues (Log)</i>				
INF_{it}	0.478 (0.304)	0.509* (0.306)	0.492 (0.301)	0.532* (0.303)
Mean	20,409,462	20,409,462	20,409,462	20,409,462
Median	213,828	213,828	213,828	213,828
N	139,600	139,600	139,600	139,600
<i>Panel B: Total Costs (Log)</i>				
INF_{it}	0.134 (0.177)	0.138 (0.178)	0.144 (0.177)	0.149 (0.178)
Mean	21,573,911	21,573,911	21,573,911	21,573,911
Median	269,888.5	269,888.5	269,888.5	269,888.5
N	139,600	139,600	139,600	139,600
<i>Panel C: Pre-Tax Profit (IHS)</i>				
INF_{it}	0.490 (0.697)	0.407 (0.684)	0.423 (0.690)	0.355 (0.677)
Mean	425,091.9	425,091.9	425,091.9	425,091.9
Median	1,050	1,050	1,050	1,050
N	139,600	139,600	139,600	139,600
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. PLACEBO TEST: Treatment cohorts randomly reshuffled among treated firms. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. *p<0.1, **p<0.05, ***p<0.01.

Table A.18: Debt and Receivables Growth (Placebo Test)

	(1)	(2)	(3)	(4)
<i>Panel A: Receivables (Log)</i>				
INF_{it}	0.164 (0.149)	0.156 (0.145)	0.158 (0.150)	0.156 (0.147)
Mean	5,246,227	5,246,227	5,246,227	5,246,227
Median	185,437	185,437	185,437	185,437
N	139,035	139,035	139,035	139,035
<i>Panel B: Short-Term Debt (Log)</i>				
INF_{it}	0.162 (0.186)	0.155 (0.184)	0.166 (0.186)	0.163 (0.184)
Mean	8,591,694	8,591,694	8,591,694	8,591,694
Median	353,976	353,976	353,976	353,976
N	139,035	139,035	139,035	139,035
<i>Panel C: Long-Term Debt (Log)</i>				
INF_{it}	0.226 (0.359)	0.219 (0.355)	0.215 (0.362)	0.213 (0.358)
Mean	2,316,866	2,316,866	2,316,866	2,316,866
Median	35,914.5	35,914.5	35,914.5	35,914.5
N	139,035	139,035	139,035	139,035
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. PLACEBO TEST: Treatment cohorts randomly reshuffled among treated firms. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. *p<0.1, **p<0.05, ***p<0.01.

Table A.19: Cost Reallocation (Placebo Test)

	(1)	(2)	(3)	(4)
<i>Panel A: Services (Log)</i>				
INF_{it}	-0.020 (0.154)	-0.004 (0.155)	-0.009 (0.153)	0.007 (0.155)
Mean	4,724,181	4,724,181	4,724,181	4,724,181
Median	69,704	69,704	69,704	69,704
N	139,600	139,600	139,600	139,600
<i>Panel B: Third-Party Assets (Log)</i>				
INF_{it}	0.082 (0.242)	0.147 (0.245)	0.077 (0.242)	0.146 (0.244)
Mean	401,009.4	401,009.4	401,009.4	401,009.4
Median	2,477	2,477	2,477	2,477
N	139,600	139,600	139,600	139,600
<i>Panel C: Labor Costs (Log)</i>				
INF_{it}	0.185 (0.244)	0.235 (0.240)	0.204 (0.244)	0.257 (0.243)
Mean	3,287,842	3,287,842	3,287,842	3,287,842
Median	17,047	17,047	17,047	17,047
N	139,600	139,600	139,600	139,600
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. PLACEBO TEST: Treatment cohorts randomly reshuffled among treated firms. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. *p<0.1, **p<0.05, ***p<0.01.

Table A.20: Claims Accumulation and Liquidity Drain (Placebo Test)

	(1)	(2)	(3)	(4)
<i>Panel A: Receivables/Assets</i>				
INF_{it}	0.000 (0.013)	-0.001 (0.013)	0.003 (0.012)	0.002 (0.013)
Mean	0.31	0.31	0.31	0.31
Median	0.24	0.24	0.24	0.24
N	138,937	138,937	138,937	138,937
<i>Panel B: Cash/Assets</i>				
INF_{it}	-0.003 (0.010)	-0.003 (0.010)	-0.005 (0.010)	-0.005 (0.010)
Mean	0.11	0.11	0.11	0.11
Median	0.03	0.03	0.03	0.03
N	139,502	139,502	139,502	139,502
<i>Panel C: Claims-Liquidity Congestion</i>				
INF_{it}	0.004 (0.018)	0.003 (0.018)	0.009 (0.018)	0.008 (0.018)
Mean	0.2	0.2	0.2	0.2
Median	0.14	0.14	0.14	0.14
N	138,937	138,937	138,937	138,937
Unit FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Province $\times$ Year FE		✓		✓
ATECO $\times$ Year FE			✓	✓

Notes: Sun & Abraham (2021) interaction-weighted estimator. INF_{it} is the aggregated ATT. PLACEBO TEST: Treatment cohorts randomly reshuffled among treated firms. Standard errors clustered at the firm level in parentheses. Mean and Median computed at $t = -1$ for treated firms. *p<0.1, **p<0.05, ***p<0.01.

Table A.21: Leave-One-Cohort-Out Robustness Check

	95% CI	LOCO Mean $\pm$ SD
Revenues	[2.499, 3.659]	3.078 $\pm$ 0.059
Total Costs	[1.427, 2.086]	1.757 $\pm$ 0.057
Pre-Tax Profit	[0.882, 2.980]	1.930 $\pm$ 0.214
Receivables	[1.588, 2.220]	1.909 $\pm$ 0.061
Short-Term Debt	[1.256, 1.913]	1.584 $\pm$ 0.058
Long-Term Debt	[1.404, 2.517]	1.982 $\pm$ 0.108
Services	[1.586, 2.334]	1.962 $\pm$ 0.042
Third-Party Assets	[1.819, 2.673]	2.256 $\pm$ 0.091
Labor Costs	[2.354, 3.266]	2.813 $\pm$ 0.094
Receivables/Assets	[0.076, 0.130]	0.104 $\pm$ 0.005
Cash/Assets	[-0.067, -0.025]	-0.046 $\pm$ 0.003
Claims-Liquidity Congestion	[0.110, 0.188]	0.149 $\pm$ 0.006

Notes: Column “95% CI” reports the 95% confidence interval of the Sun & Abraham (2021) ATT estimate using all treatment cohorts. Column “LOCO Mean $\pm$ SD” reports the mean and standard deviation of ATT estimates from leave-one-cohort-out analysis, where each estimate excludes one treatment cohort. Estimates are from log-transformed outcomes (IHS for profit). Standard errors clustered at the firm level.

B. Temporal Graph Network

This appendix provides a formal and self-contained description of the Temporal Graph Network (TGN) architecture adopted in this study, including detailed formulations of all components summarized in the main text. Let $\mathcal{G}_t = (V, E_t)$ denote a temporal directed graph where E_t is the set of events observed up to time t . Each event is a quadruple $(u, v, t, \mathbf{x}_{uv}^t)$ representing an interaction between source node u and target node v at time t with associated edge features $\mathbf{x}_{uv}^t$.

Event Routing and Message Construction

Each temporal event is a quadruple $(u, v, t, \mathbf{x}_{uv}^t)$. The model first performs *event routing*: if $u = v$ the event is treated as a *self-loop update* and no graph message passing or scoring is performed; if $u \neq v$ the full TGN pipeline is applied.

When an ownership event occurs between two nodes u and v at time t , the model generates two directional messages ($u \rightarrow v$ and $v \rightarrow u$). Each message aggregates the current memories of source and destination nodes, relevant edge attributes (e.g., ownership share or change), and the elapsed time since their last interaction. Temporal information enters the model through a *relative time encoder*, which maps time differences into continuous vectors that modulate the attention weights. These messages are routed by a dedicated *Block Message Module*, which handles edge events (ownership creations, updates, removals) and self-loop events (firm-level attribute updates) through distinct processing paths. This routing mechanism is described in detail in the Block Message Module subsection below.

Memory Read For any node $i \in \{u, v\}$, let

$$\mathbf{h}_i^{t-} = \text{Memory}[i]$$

denote its state immediately prior to event time t .

Case $u \neq v$: Edge Events with Attention

Historical Neighbor Construction For each endpoint $i \in \{u, v\}$ the model retrieves the most recent K temporal neighbors $\mathcal{N}_K(i, t)$. For a past interaction (i, j, t_j) with

$t_j < t$, the neighbor message is

$$\mathbf{m}_{i,j}^{\text{hist}} = \left[\mathbf{x}_{ij}^{t_j} \parallel \phi(t - t_j) \parallel \mathbf{h}_j^{t_j^-} \right],$$

where ϕ denotes the temporal encoding. A permutation-invariant aggregator produces

$$\mathbf{M}_i^{\text{hist}} = \text{Agg}\{\mathbf{m}_{i,j}^{\text{hist}} : j \in \mathcal{N}_K(i, t)\}.$$

Attention Mechanism The attention module is implemented using the **TransformerConv** operator from PyTorch Geometric (Fey and Lenssen, 2019), which provides a graph-based variant of query-key-value attention (Vaswani et al., 2017).¹¹

For node u , the input representation is

$$\mathbf{s}_u^t = \left[\mathbf{h}_u^{t^-} \parallel \mathbf{M}_u^{\text{hist}} \right],$$

while each neighbor $i \in \mathcal{N}_K(u, t)$ contributes

$$\mathbf{s}_{u,i}^{t_i} = \left[\mathbf{h}_i^{t_i^-} \parallel \mathbf{m}_{u,i}^{\text{hist}} \right].$$

A single-head attention mechanism is used. Queries, keys, and values are computed as

$$\mathbf{q}_u = W_Q \mathbf{s}_u^t, \quad \mathbf{k}_{u,i} = W_{\text{neigh}} \mathbf{s}_{u,i}^{t_i}, \quad \mathbf{v}_{u,i} = W_V \mathbf{s}_{u,i}^{t_i},$$

with W_Q , W_{neigh} , W_V learnable matrices.

Attention weights are

$$\alpha_{u,i} = \frac{\exp(\mathbf{q}_u^\top \mathbf{k}_{u,i} / \sqrt{d})}{\sum_{j \in \mathcal{N}_K(u, t)} \exp(\mathbf{q}_u^\top \mathbf{k}_{u,j} / \sqrt{d})}.$$

¹¹Unlike the original Transformer self-attention mechanism, **TransformerConv** performs attention over a node’s graph neighborhood rather than over a full sequence and omits components such as multi-layer feed-forward blocks and positional encodings. It can therefore be viewed as a localized and structurally constrained adaptation of Transformer-style attention to graph data.

The neighbor context vector is

$$\mathbf{c}_u = \sum_{i \in \mathcal{N}_K(u,t)} \alpha_{u,i} \mathbf{v}_{u,i}.$$

The embedding of node u at time t combines a self-term and an attention-based neighbor term:

$$\mathbf{z}_u^t = W_{\text{self}} \mathbf{m}_u^{t-} + W_{\text{neigh}} \mathbf{c}_u.$$

The embedding $\mathbf{z}_v^t$ is computed analogously.

Scoring Architecture and Link Prediction Loss

Link prediction uses a single-tower scoring architecture, in which the ordered concatenation

$$[\mathbf{z}_u^t \| \mathbf{z}_v^t]$$

is passed through a Multi-Layer Perceptron (MLP)¹² to produce

$$\hat{y}_{uv}^t = \sigma(\text{MLP}([\mathbf{z}_u^t \| \mathbf{z}_v^t])).$$

Because the MLP receives an *ordered* pair, the resulting score is directional, i.e. $\hat{y}_{uv}^t \neq \hat{y}_{vu}^t$ in general. The same scoring architecture is used in the full-data and in the first-year variant.

The link prediction loss is the binary cross-entropy with negative sampling,

$$\mathcal{L}_{\text{link}} = -\log \hat{y}_{uv}^t - \sum_{(u,n) \in \mathcal{N}^{\text{neg}}} \log(1 - \hat{y}_{un}^t).$$

For each observed (positive) edge (u, v, t) , ten negative samples are generated by pairing the same source node u with randomly selected destination nodes v' that have never been connected to u up to time t . The negative sampler uses a “filtered” strategy, ensuring that no existing edge is inadvertently used as a negative example.

¹²An MLP, or Multi-Layer Perceptron, is a type of artificial neural network with at least three layers: an input layer, one or more hidden layers, and an output layer.

Regression Heads

Beyond link prediction, the model also produces regression outputs by attaching separate regression heads $g(\cdot)$ to the node embeddings. Two regression tasks are defined:

- **Level Regression Head:** predicts the current ownership share (s_{uv}) in log scale for active links, optimized through an ℓ_1 loss for robustness to moderate noise.
- **Delta Regression Head:** predicts the change in log-ownership share (Δs_{uv}) between consecutive observations, optimized via a Huber loss to handle heavy-tailed variations.

For each active link with regression target y_i^t , the model predicts

$$\hat{y}_i^t = g(\mathbf{z}_i^t),$$

and the corresponding regression loss is

$$\mathcal{L}_{\text{reg}} = \frac{1}{|\mathcal{I}_t|} \sum_{i \in \mathcal{I}_t} \ell(\hat{y}_i^t, y_i^t),$$

where ℓ is either the mean absolute error (MAE) or the squared error used to compute RMSE, depending on the task. Regression heads share the same embeddings as the link prediction component but do not alter the event memory update mechanism.

Dynamic Loss Weighting

The three prediction heads share the same encoder but contribute differently depending on the event type: creation events primarily activate the link head; updates trigger both regression heads; and removals are supervised through the delta head. Their contributions to the total objective are balanced dynamically through two convex coefficients, γ and η , which respectively regulate the trade-off between link versus regression losses, and between level versus delta regression components. In order to balance the three tasks during training, η is set to 0.5, while γ , i.e. the importance of the regression task relative to link prediction, is defined dynamically for each batch as follows:

$$\gamma = \min \left(\frac{n_{\text{upd}} + n_{\text{del}}}{n_{\text{upd}} + n_{\text{del}} + n_{\text{cre}}}, 0.5 \right),$$

where n_{cre} , n_{upd} , and n_{del} denote respectively the number of *creation*, *update*, and *deletion* events contained in the current mini-batch. Intuitively, γ increases when the batch contains a larger proportion of non-creation events, encouraging the model to allocate greater weight to regression tasks when ownership relations are being actively modified or terminated. Conversely, when new link formations dominate, γ decreases, shifting the focus toward learning link existence patterns. The upper cap of 0.5 ensures that the link-prediction component always retains sufficient gradient signal, preventing the model from collapsing into a purely regression-driven regime during periods of intense update activity.

Auxiliary Losses

Barlow Twins Loss To further regularize the embeddings, we include a Barlow Twins-style redundancy-reduction term (Zbontar et al., 2021). For each node i , we generate two stochastic views of its embedding, $\mathbf{z}_i^{(1),t}$ and $\mathbf{z}_i^{(2),t}$ (e.g. via dropout¹³), and form the empirical cross-correlation matrix $C \in \mathbb{R}^{d \times d}$ with entries

$$C_{ab} = \frac{1}{|\mathcal{I}_t|} \sum_{i \in \mathcal{I}_t} z_{i,a}^{(1),t} z_{i,b}^{(2),t}.$$

The Barlow Twins loss is

$$\mathcal{L}_{\text{BT}} = \sum_a (1 - C_{aa})^2 + \lambda_{\text{off}} \sum_{a \neq b} C_{ab}^2,$$

which enforces invariance across views on the diagonal and discourages redundancy across dimensions on the off-diagonal.

Temporal Consistency Loss We impose a temporal smoothness constraint on node embeddings to penalize abrupt changes over time (Trivedi et al., 2019). For nodes observed at consecutive instants, the temporal loss is

$$\mathcal{L}_{\text{temp}} = \frac{1}{|\mathcal{I}_t|} \sum_{i \in \mathcal{I}_t} \left\| \mathbf{z}_i^t - \mathbf{z}_i^{t-} \right\|_2^2,$$

¹³Dropout randomly zeroes out a subset of embedding dimensions, producing two stochastically different *views* of the same embedding vector.

where $\mathbf{z}_i^{t-}$ denotes the embedding of node i immediately prior to time t (when defined).

Total Loss

The resulting total loss function is defined as a convex combination of the link prediction and regression components, augmented with the two auxiliary regularization terms:

$$\mathcal{L}_{\text{total}} = (1 - \gamma) \mathcal{L}_{\text{link}} + \gamma [(1 - \eta) \mathcal{L}_{\text{level}} + \eta \mathcal{L}_{\Delta}] + \lambda_{\text{NC}} \mathcal{L}_{\text{Barlow}} + \lambda_{\text{TMP}} \mathcal{L}_{\text{temporal}}.$$

Here, $\mathcal{L}_{\text{link}}$ is the binary cross-entropy loss for link prediction, $\mathcal{L}_{\text{level}}$ and $\mathcal{L}_{\Delta}$ are respectively the ℓ_1 and Huber regression losses for the level and delta tasks, while $\mathcal{L}_{\text{Barlow}}$ and $\mathcal{L}_{\text{temporal}}$ are auxiliary terms enforcing, respectively, representation decorrelation and temporal smoothness.¹⁴ This formulation enables the network to jointly learn the structure, intensity, and evolution of ownership relations while maintaining stable and temporally consistent embeddings.

Memory Update

Each endpoint receives an event-level message, e.g.

$$\mathbf{m}_u^t = [\mathbf{x}_{uv}^t \parallel \phi(t - t_u^{\text{last}}) \parallel \mathbf{z}_u^t],$$

and the memory is updated via

$$\mathbf{h}_u^t = \text{GRU}(\mathbf{h}_u^{t-}, \mathbf{m}_u^t), \quad \mathbf{h}_v^t = \text{GRU}(\mathbf{h}_v^{t-}, \mathbf{m}_v^t).$$

Case $u = v$: Self-loop Events

Self-loop events update the memory of the involved node but do not trigger message passing, attention, scoring, or loss computation. For such events, the model forms a simplified update message $\mathbf{m}_u^t$ based solely on the event features and temporal encoding, and applies

$$\mathbf{h}_u^t = \text{GRU}(\mathbf{h}_u^{t-}, \mathbf{m}_u^t).$$

¹⁴The weights of these two losses are $\lambda_{\text{NC}} = 0.1$ and $\lambda_{\text{TMP}} = 0.01$, respectively.

No embedding $\mathbf{z}_u^t$ is generated for prediction, and the event contributes no gradient to the scoring model.

Block Message Module

The Block Message Module handles the routing between edge events and self-loop events through distinct processing paths. For edge events ($u \neq v$), the time-encoded messages are concatenated with edge-level features and passed to the attention-based aggregation mechanism described above. For each such event, two directional messages are constructed ($u \rightarrow v$ and $v \rightarrow u$), each aggregating the current memories of source and destination nodes, the relevant edge attributes (e.g., ownership share or its change), and the elapsed time since the last interaction between the pair. Temporal information enters the model through a *relative time encoder*, which maps time differences into continuous vectors that modulate the attention weights.

In contrast, self-loop events ($u = v$)—which correspond to yearly firm-level attribute updates—bypass the attention layer and directly update the node’s memory, ensuring that firm attributes remain temporally aligned with the network’s evolution without generating spurious interaction signals. These are yearly updates applied to firm nodes (PG) where $u = v$, used to refresh firm-level attributes such as financial indicators and judicial status. Firms that are confiscated undergo a discrete feature change at the moment of seizure, when their status variable switches from “not confiscated” to the relevant judicial stage, and their “confiscated share” variable updates from 0% to the seized percentage. Individuals (PFs), lacking such information, never generate self-loop events and evolve purely through their relational activity. This asymmetry is handled with a firm-level mask ensuring that node features are updated only for firms. In the first-year variant, the only self-loop event a firm generates is the initial one, stamped just before its first ownership event and carrying no judicial information; all of its later memory updates come from edge events.

This design preserves the separation between structural changes (edges) and internal updates (nodes), while allowing both to contribute to the latent dynamics of ownership relations.

Training Performance and Diagnostics

Table A.4 compares the two variants on the held-out test events. Both achieve strong link-prediction performance, with an AUC of 0.98 for the full-data model and 0.96 for the first-year model, and average precision of 0.83 and 0.74 respectively, indicating that the TGN captures the temporal and structural patterns of link formation in the ownership network even when firm attributes are observed only once. The gap between the two is expected: the first-year model must predict ownership events without ever seeing a firm’s current balance sheet, and it is trained on a much shorter event stream. For the regression tasks, mean absolute errors of about 0.71 and 0.77 on log-transformed stake levels are moderate relative to the dispersion of the target ($\text{IQR} = 3.38$), given the skewness of stake magnitudes—from minority holdings to full control. Performance on the delta task is similar across variants, with $\text{MAE}_\Delta \approx 0.72$, and inherently noisy: nearly 70% of stake changes are zero, while large negative deltas (full or partial divestments) dominate the non-zero tail. Figure B.15 shows that the training loss of the full-data model declines smoothly over the 20 epochs, with little dispersion across seeds.

We emphasize that these predictive metrics serve only as indirect indicators of representation quality. The purpose of this evaluation is not to deploy the model as an operational predictor. Instead, our goal is to verify that the learned embeddings capture the latent structure and temporal dependencies of the ownership network. Given that, the results suggest the model has learned meaningful representations. These diagnostics establish that the model learns the generative structure of ownership dynamics; the main text describes how embeddings are extracted without leakage before constructing the proximity-based risk index used in the validation and empirical sections.

Hyperparameters and Optimization

Unless otherwise stated, all experiments use the following hyperparameter configuration for the TGN model and training procedure.

Model architecture

- Memory / embedding dimension: $d = 64$.
- Time encoding dimension: $d_{\text{time}} = 100$.

- Temporal neighbors per node: $K = 50$ (most recent neighbors).
- GNN layer: **TransformerConv** with
 - number of heads: $H = 2$,
 - dropout rate: 0.2,
 - input and output channels both equal to the memory dimension.
- Negative sampling: “filtered” sampler with 10 negatives per positive edge.


Training and optimization.

- Optimizer: Adam with learning rate $\eta_0 = 10^{-3}$ and weight decay 10^{-5} .
- Batch size: 256 events.
- Number of epochs per model: 20.
- Gradient clipping: global ℓ_2 -norm clipped at 1.0.
- Learning rate scheduler: **ReduceLROnPlateau** on validation AP.
- Regularization weights:
 - node-level Barlow Twins loss: $\lambda_{\text{node-contrast}} = 0.10$,
 - temporal consistency loss: $\lambda_{\text{temporal}} = 0.05$.

Evaluation metrics. Link prediction performance is assessed using the Area Under the ROC Curve (AUC) and Average Precision (AP), while the regression heads are evaluated with Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). All metrics are computed in a strictly time-respecting fashion, with test events occurring after all training and validation intervals. After convergence, we extract the time-indexed embeddings $z_t(\cdot)$ from the final memory state, obtaining compact representations that summarize each firm’s evolving relational behavior and exposure within the ownership network. The extracted embeddings have 64 dimensions.

Data splits. Model parameters are updated after each batch, and node memories are carried forward across batches to preserve temporal continuity. In the full-data variant,

Figure B.15: Training Loss of the Full-Data TGN




Notes: Training loss of the full-data TGN over 20 epochs. The solid line is the average across three random seeds, the band ± 1 standard deviation across seeds.

firm-level attributes are refreshed every year via self-loop (PG) events, ensuring that structural and attribute dynamics remain synchronized. Evaluation follows the same event-driven procedure: the chronological data split assigns events from 1993 to end of 2015 for training, end of 2015 to early 2019 for validation, and 2019 to 2024 for testing. Figure A.3 breaks down the events by the data splits.

The model is implemented using PyTorch (Paszke et al., 2019) and PyTorch Geometric (Fey and Lenssen, 2019), and the architecture comprises roughly 206k trainable parameters in both variants. Each variant is trained with three random seeds (42, 123 and 456), and the embeddings used downstream are those of the best-performing seed. Figure B.16 illustrates the overall TGN architecture, while Figure B.15 reports the training loss of the full-data model over the 20 epochs.

Figure B.16: Temporal Graph Network Architecture



Notes: The figure illustrates the architecture of the Temporal Graph Network (TGN) used to generate firm-level embeddings from the dynamic ownership network. The model processes a stream of time-stamped events, computing context-aware embeddings through message passing and attention mechanisms and updating node memories. Node events only contribute to the memory update

C. Additional Embedding Validation

This appendix collects supplementary evidence on the properties of the ownership embeddings learned by the Temporal Graph Network, complementing the main results reported in Section 5. It reports a low-dimensional visualization of the embedding space, an alternative distance-based proximity measure and its distribution across snapshot years, and the full comparison between embedding neighborhoods and observed graph neighborhoods. The purpose is strictly validation-oriented, and no causal interpretation is attached to the analyses presented here. Unless stated otherwise, all exercises use the full-data representation.

Visualizing the Embedding Space Figure C.17 illustrates the space through a t-SNE projection of the 64-dimensional embeddings in 2006 and 2010. Individuals (*Persone Fisiche*, white circles) form a cluster distinct from firms, indicating that the model separates node types from relational and attribute event patterns alone, and within the firm mass confiscated firms (red) appear in localized sub-clusters rather than mixing uniformly with the rest (green). The plots are a visual aid only: t-SNE is a non-linear, stochastic reduction whose distances and cluster shapes are not faithful to the geometry of the original space, so it cannot establish separation, proximity, or risk. All substantive claims rest on the quantitative exercises of Section 5, conducted in the original embedding space.


An Alternative Proximity Measure Section 5.1 summarizes a firm’s position through its projection on the confiscation direction and through the kernel-weighted IPI. Here we report a third, distance-based summary, which shares the time-respecting reference set of the IPI but has the advantage of an interpretable distance scale. We ask whether firms that will be confiscated in the future already occupy regions of the embedding space closer to firms confiscated in the past, at multiple time snapshots.

For each year t , we define the reference set

$$\mathcal{C}_t = \{j : \text{firm } j \text{ is confiscated at or before } t\},$$

which contains all firms that have experienced judicial confiscation up to that point. We then consider all remaining firms $i \notin \mathcal{C}_t$, including both firms that will be confiscated at

Figure C.17: Embeddings Visualization (2006 and 2010)



Notes: Embeddings for different years capture the evolving relationships and structures within the ownership network. Red nodes represent confiscated firms, yellow nodes not-yet-confiscated firms, green nodes represent non-confiscated firms and white nodes people. UMAP is used for dimensionality reduction in this 2D visualization. Full-data representation.

a later date and firms that are never confiscated.

Let z_i and z_j denote the learned embedding vectors of firms i and j . We measure proximity in the latent space by computing cosine distances

$$d_{ij} = 1 - \frac{z_i \cdot z_j}{\|z_i\| \|z_j\|},$$

where smaller values indicate greater similarity along latent relational dimensions. To aggregate distances to all firms in the confiscated reference set into a smooth and robust proximity measure, we apply an exponential kernel,

$$w_{ij} = \exp(-\gamma d_{ij}),$$

with bandwidth parameter $\gamma > 0$. The resulting similarity score for firm i at time t is defined as

$$S_i(t) = \frac{1}{|\mathcal{C}_t|} \sum_{j \in \mathcal{C}_t} w_{ij}.$$

For interpretability, we map this quantity into a distance-to-cluster measure,

$$d_i(t) = -\frac{1}{\gamma} \log S_i(t),$$

where lower values indicate closer proximity to the region of the embedding space populated by confiscated firms.

We compute these proximity measures for each firm across multiple snapshot years and compare their distributions between firms that are never confiscated and firms that will be confiscated at some future date. Figure C.18 reports the resulting distributions for the years 2000 and 2010. Across both snapshots, firms that will be confiscated in the future exhibit a pronounced leftward shift, concentrating at substantially lower distance values relative to firms that are never confiscated. This pattern is the distance-based counterpart of the result reported in Figure 2: future-confiscated firms already occupy regions of the latent space closer to the contemporaneously known confiscated set well before any judicial intervention occurs. Because $\mathcal{C}_t$ contains only firms confiscated at or before t , and because future confiscation enters solely as an ex-post grouping label, the comparison does not use information unavailable at the snapshot date.

Firms that are never confiscated display larger distances on average, reflecting weaker similarity to entities whose confiscation is already known. At the same time, their distribution is heterogeneous, with a non-negligible mass at relatively low distance values: a subset of firms that are never subject to confiscation nonetheless share relational characteristics with confiscated entities. This overlap should temper any reading of the score as a classifier; it separates the groups on average without cleanly partitioning them.

Graph Topology and Latent Structure We next assess whether the embeddings merely reproduce observed network connectivity or instead capture higher-order and non-local relational information. To this end, we examine the overlap between neighborhoods defined in the embedding space and one-hop neighborhoods in the observed ownership graph.

For each firm i , let $N_G(i)$ denote its set of one-hop neighbors in the network and $N_k(i)$ its k nearest neighbors in the embedding space. We evaluate neighborhood alignment

using Precision@ k and Recall@ k :

$$\text{Precision@}k = \frac{|N_G(i) \cap N_k(i)|}{k}, \quad \text{Recall@}k = \frac{|N_G(i) \cap N_k(i)|}{|N_G(i)|}.$$


Table C.22 reports average Precision@ k and Recall@ k for $k \in \{5, 10, 20\}$ in five snapshot years, together with the expected precision under random selection. Two features stand out. Embedding neighborhoods recover graph neighbors far more often than chance: Precision@5 ranges from about 0.010 to 0.032 against a random benchmark of roughly 0.00002, a ratio of several hundred to one, in a network whose average degree is close to unity. At the same time, the absolute overlap is small in both directions: Precision@ k never exceeds about 3%, and Recall@20 peaks at roughly 17%. The right reading is therefore that embedding proximity is related to observed topology without mechanically reproducing one-hop adjacency. Most of a firm’s nearest neighbors in the latent space are not its direct counterparties, and most of its direct counterparties are not among its nearest latent neighbors; the representation aggregates relational history and higher-order structure rather than restating the adjacency matrix.

Table C.22: KNN–graph neighbourhood overlap by year


	2000	2005	2010	2015	2020
Panel A: $k = 5$					
Precision@5	0.0318	0.0122	0.0096	0.0120	0.0117
Recall@5	0.0850	0.0304	0.0249	0.0315	0.0341
Baseline precision	0.00002	0.00002	0.00002	0.00001	0.00001
Panel B: $k = 10$					
Precision@10	0.0223	0.0093	0.0072	0.0086	0.0093
Recall@10	0.1160	0.0467	0.0397	0.0469	0.0522
Baseline precision	0.00002	0.00002	0.00002	0.00001	0.00001
Panel C: $k = 20$					
Precision@20	0.0164	0.0071	0.0054	0.0067	0.0077
Recall@20	0.1725	0.0732	0.0599	0.0726	0.0824
Baseline precision	0.00002	0.00002	0.00002	0.00001	0.00001
Average degree	1.13	1.23	1.24	1.18	1.12

Note: This table presents the overlap between k -nearest neighbours in embedding space and 1-hop graph neighbourhoods across different years. Precision@ k measures the fraction of embedding KNNs that are true graph neighbours, while Recall@ k measures the fraction of graph neighbours recovered by the embedding KNN. The baseline precision represents the expected precision from randomly selecting k neighbours.

Figure C.18: Distance to the Contemporaneously Confiscated Set, by Future Confiscation Status



(a) Snapshot year: 2000



(b) Snapshot year: 2010

Notes: Distribution of the distance-to-cluster measure $d_i(t)$ defined above, for two snapshot years. The unit of observation is a firm present in that year and not already confiscated. Lower values indicate closer proximity to the reference set $\mathcal{C}_t$ of firms confiscated at or before the snapshot year, so no later seizure enters the measure. Green: never confiscated. Red: confiscated after the snapshot year, a label used only to split the sample. Full-data representation.

D. Forecasting Confiscation: Design and Additional Results

This appendix provides a self-contained account of the forecasting exercise summarized in Section 5.2. It documents the prediction target and horizon convention, the construction of the classification sample, the two information sets, the preprocessing and rebalancing steps, the chronological evaluation design, the classifiers and their tuning, and the complete set of results. A reader should be able to reconstruct the design from this appendix alone.

D.1. Prediction Target and Forecast Horizons

The object of interest is the probability p_i that firm i is confiscated, given the information $\mathbf{X}_i$ available at the time of the forecast:

$$p_i = f(\mathbf{X}_i).$$

Confiscation is dated at the initial judicial seizure of assets (*sequestro*), as everywhere else in the paper, and is denoted t_i^c for firm i .

Every observation is aligned to a *reference year*. For a firm and a horizon h , the reference year is

$$t_{i,h}^c = t_i^c - h,$$

and the feature vector is built from information dated at the reference year and at the years preceding it. Because the most recent information entering the classifier is measured at $t_i^c - h$ while confiscation occurs at t_i^c , the horizon h is the number of years by which the forecast anticipates the judicial event: $h = 1$ is a one-year-ahead forecast and $h = 4$ a four-year-ahead forecast. We estimate the exercise separately for $h = 1, 2, 3, 4$. This convention is applied uniformly to the equations, the prose, the tables and the figures of this appendix and of Section 5.2.

The eligible population at each horizon consists of firms (*Persone Giuridiche*) present in the yearly ownership graphs with at least the number of consecutive prior observations required by the lag structure. Individuals (*Persone Fisiche*) are excluded throughout,

since they cannot be confiscated entities. Firms without sufficient history at a given horizon are dropped, which is why the sample shrinks as h increases.

D.2. Classification Sample and Reference Years

We construct a firm-year panel from the dynamic ownership graphs and merge onto it the firm-level variables of Table A.1 and the yearly Infiltration Proximity Index of Section 5.1.

For confiscated firms the reference year is observed directly. Non-confiscated firms have no confiscation date, and simply assigning them an arbitrary year would confound the comparison with calendar-time composition: confiscations are concentrated in particular years (Figure A.1), so a naive assignment would let the classifier separate the two groups on the basis of when they are observed rather than how they are positioned. We therefore assign each non-confiscated firm a *synthetic* reference year through propensity-score matching, pairing it with the confiscated firm it most closely resembles on stable attributes—industry at the NACE two-digit level, size class, and location—and transferring that firm’s confiscation year. Figure D.19 compares the resulting distribution of synthetic reference years for non-confiscated firms with the observed distribution of confiscation years, showing that the two align closely.

Table D.23 reports the number of observations and the share of positive outcomes at each horizon. The sample declines from 13,182 observations at $h = 1$ to 10,675 at $h = 4$, reflecting the history requirement, while the share of confiscated firms falls from 17.6% to 12.8%.

Table D.23: Descriptive Statistics by Forecast Horizon


	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Observations	13,182	12,007	11,611	10,675
Mean Confiscation	0.1761	0.1554	0.1422	0.1284

Notes: The table reports descriptive statistics for the classification dataset across different forecast horizons. $h > 0$ indicates the number of years ahead for the prediction task. The number of observations decreases for higher forecast horizons due to data availability constraints.

D.3. Information Sets

Two information sets are compared, each estimated separately at every horizon.

Figure D.19: Matching Strategy for Synthetic Reference Year Assignment



Notes: Reference years used to align observations in the classification sample. Blue: observed confiscation years. Green: synthetic reference years assigned to non-confiscated firms by propensity-score matching on industry (NACE two-digit), size class and location. Aligning the groups in calendar time prevents the classifier from separating them on when they are observed.

(i) *Firm-level data.* Firm characteristics and financial indicators drawn from the business registry and from balance sheets: numeric variables $\mathbf{x}_{i,t}$ and categorical variables $\mathbf{c}_{i,t}$. Table A.1 lists them in full, grouped into firm characteristics, financial indicators, balance-sheet assets, balance-sheet liabilities and equity, and income-statement items. All continuous variables enter with five lags ($L = 5$) and all categorical variables with two ($K = 2$), in addition to their value at the reference year, for roughly 560 features in all.

(ii) *IPI only.* The Infiltration Proximity Index $\text{IPI}_{i,t}$ of Section 5.1, at the reference year and at four lags ($L_{\text{IPI}} = 4$), five features in all. The index is computed on the full-data representation of Section 4, with the time-respecting reference set: in year t only firms confiscated at or before t count as confiscated. Lags are taken within firm on the yearly panel of the index. We stress a point of interpretation. Because the full-data TGN receives yearly firm information through financial self-loop events, this information set is *not* free of firm-level financial information: part of it may already be

encoded in the embeddings on which the index is built. Set (ii) should be understood as a one-dimensional summary of the ownership space that jointly reflects the ownership and attribute history observed by the model, and the comparison with set (i) as a comparison between that summary and the same underlying variables entered directly, with their full lag structure. Rerunning the exercise with the index computed on the first-year representation would separate the two channels; we regard this as the natural extension of the exercise.

Writing $t_h \equiv t_{i,h}^c$ for compactness, the two feature vectors are

$$\begin{aligned}\mathbf{X}_{i,h}^{\text{firm}} &= \left[\mathbf{x}_{i,t_h}, \dots, \mathbf{x}_{i,t_h-L}, \mathbf{c}_{i,t_h}, \dots, \mathbf{c}_{i,t_h-K} \right], \\ \mathbf{X}_{i,h}^{\text{IPI}} &= \left[\text{IPI}_{i,t_h}, \dots, \text{IPI}_{i,t_h-L_{\text{IPI}}} \right].\end{aligned}$$

Each observation is labelled with a binary target $y_{i,h}$ equal to one if the firm is confiscated h years after its reference year and zero otherwise.

A timing detail matters for the index. Yearly embedding snapshots are extracted at January 1 of each year, so the index labelled year T reflects only ownership events processed before that date and only seizures recorded before it (Section 4). Balance-sheet variables labelled year T , by contrast, refer to the accounting year ending on December 31 of T . The IPI block is therefore, if anything, more conservatively dated than the firm-level block.

D.4. Preprocessing and Class Imbalance

Missing values in continuous variables are linearly imputed within firm where a series is observed on both sides of the gap, and flagged otherwise; numerical variables are then standardized. Categorical variables are retained as such and handled natively by the tree-based learners, with one-hot encoding applied for the linear classifier.

Confiscated firms are a minority of the sample at every horizon (Table D.23). To prevent the classifiers from collapsing onto the majority class, we apply the Synthetic Minority Over-sampling Technique (Chawla et al., 2002) to the training data only, inside each cross-validation fold, so that no synthetic observation can influence a validation or test score. SMOTE creates new minority observations rather than duplicating existing ones: for a minority instance $\vec{x}_i$, one of its k nearest same-class neighbors $\vec{x}_{zi}$ is drawn

at random and a synthetic point is placed on the segment joining them,

$$\vec{x}_{new} = \vec{x}_i + \lambda(\vec{x}_{zi} - \vec{x}_i), \quad \lambda \sim U[0, 1],$$

which populates the feature space around minority observations and yields a less sparse decision region for that class. Test sets are never rebalanced, so reported out-of-sample metrics refer to the empirical class distribution.

D.5. Chronological Evaluation Design

Two distinct chronological cut-offs discipline the exercise, and they should not be conflated.

TGN parameter estimation. The temporal graph network is trained on the event stream running from 1993 to the end of 2015, with the window from the end of 2015 to early 2019 used for validation monitoring and the remainder held out. Its parameters are therefore estimated on pre-2016 events. Events occurring later are streamed through the trained model in evaluation mode: node memories are updated sequentially as the events arrive, but no gradient is taken and no parameter is modified. This is deliberate, and it defines the sense in which later embeddings are “out of sample”. They are out of sample with respect to *parameter estimation*; they are not generated without observing later ownership events, since the model does observe them as they arrive. This mirrors how a deployed system would operate—a trained model processing new ownership events without retraining—and it is why the embeddings for later years can be used at all.

Downstream classifier. For each horizon h , the classifier is trained on observations whose reference year falls at or before 2016 $- h$ and evaluated on observations with strictly later reference years. Since the reference year is $t_i^c - h$, this is equivalent to training on firms whose (observed or synthetic) confiscation year is at or before 2016 and testing on firms confiscated afterwards. Table D.24 summarizes the resulting date ranges. Within the training window, hyperparameters are selected by five-fold time-series cross-validation, in which each fold’s training portion always precedes its validation portion, so that no fold uses future information to predict the past.

Table D.24: Chronological Design of the Forecasting Exercise

<i>Panel A: TGN parameter estimation (common to all horizons)</i>				
Training events	1993 – 31 Dec 2015			
Validation events	31 Dec 2015 – early 2019			
Later events	evaluation mode only: memory updates, no gradient			

<i>Panel B: Downstream classifier</i>				
	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Reference year, training	≤ 2015	≤ 2014	≤ 2013	≤ 2012
Reference year, test	> 2015	> 2014	> 2013	> 2012
Confiscation year, training	≤ 2016	≤ 2016	≤ 2016	≤ 2016
Confiscation year, test	> 2016	> 2016	> 2016	> 2016

Notes: Panel A refers to the temporal graph network; Panel B to the supervised classifier that takes its output as an input. The reference year is $t_{i,h}^c = t_i^c - h$, with t_i^c observed for confiscated firms and synthetic otherwise (Section D.2). The cut-off $2016 - h$ on the reference year is equivalent to a common 2016 cut-off on the confiscation year, so the training window contains the same firms at every horizon. Hyperparameters are chosen by five-fold time-series cross-validation within the training window, and SMOTE is applied inside training folds only.

D.6. Classifiers and Hyperparameter Selection

Following [Ash et al. \(2025\)](#), we compare four classifiers of increasing flexibility.

Logistic Regression A linear benchmark. Its role is diagnostic as much as competitive: comparing it with the ensembles indicates how much of the available signal is linearly accessible in a given feature space.

Random Forest An ensemble of decision trees ([Breiman, 2001](#)), whose prediction is the majority vote across trees. Two devices control the variance of individual trees: each tree is fit on a bootstrap sample of the training data, and each split considers only a random subset of features, which decorrelates the trees.

LightGBM A gradient-boosting framework ([Ke et al., 2017](#)) that builds trees sequentially, each correcting the residual errors of its predecessors. Its efficiency rests on gradient-based one-side sampling, which retains all poorly fit observations while subsampling well-fit ones when computing gradients, and on exclusive feature bundling, which merges features that are rarely non-zero simultaneously.

XGBoost A second gradient-boosting implementation ([Chen and Guestrin, 2016](#)),

which adds explicit ℓ_1 and ℓ_2 penalties on model complexity to the objective, learns default directions for missing values at each node, and parallelizes tree construction.

Hyperparameters are selected by randomized search with ten draws per model-feature combination and five-fold time-series cross-validation, scoring candidates by ROC-AUC. Randomized search samples configurations from the specified ranges rather than evaluating the full grid, which is computationally far cheaper and typically comparable in performance. The search spaces are:

- **Random Forest:** `n_estimators` $\in \{100, 200, 300\}$; `max_depth` $\in \{10, 20, 30, \text{None}\}$; `min_samples_split` $\in \{2, 5, 10\}$.
- **LightGBM:** `n_estimators` $\in \{100, 200, 400\}$; `learning_rate` $\in \{0.01, 0.1, 0.2\}$; `num_leaves` $\in \{20, 31, 40\}$.
- **XGBoost:** `n_estimators` $\in \{100, 200, 300\}$; `learning_rate` $\in \{0.01, 0.1, 0.2\}$; `max_depth` $\in \{3, 5, 7\}$.

The selected configuration is refit on the full training window and applied once to the holdout period. Reported metrics are accuracy, precision, recall, F_1 and ROC-AUC, all computed out of sample. Given the class imbalance, ROC-AUC and the precision-recall pair are the informative statistics; accuracy is reported for completeness but is mechanically high whenever the majority class dominates. Models are implemented in `scikit-learn`, with SMOTE from `imblearn` embedded in the estimation pipeline so that rebalancing is refit within each fold.


D.7. Complete Forecasting Results

Figure D.20 reports the ROC curves for the two shorter horizons, complementing the longer horizons shown in Figure 3 of the main text. Tables D.25 to D.28 report every metric for every classifier and information set, one table per horizon; decisions are taken at the classifiers' 0.5 threshold, and bold entries mark the better information set for a given classifier and metric.

The pattern is uniform across the four horizons. Firm-level data yield ROC-AUC between 0.63 and 0.75: within a confiscation-centered network, standard financial indicators discriminate weakly. The index alone reaches between 0.83 and 0.90, above the

firm-level block in all sixteen horizon–classifier combinations, and the spread across classifiers is narrow—logistic regression attains 0.895 at $h = 1$ and 0.860 at $h = 4$, against 0.903 and 0.878 for the best boosting model—indicating that the signal is close to linear in the index and its lags. The two sets differ most in the precision–recall trade-off. Firm-level data paired with the tree-based learners are the more precise option (precision between 0.54 and 0.74) but recover only 0.26–0.41 of the confiscated firms; logistic regression on firm-level data trades the other way, with recall of 0.55–0.62 at precision of 0.36–0.44. The index recovers between 0.76 and 0.92 of confiscated firms in every combination, at precision between 0.46 and 0.61, against a base rate of 13–18%. Performance does not decline as the horizon lengthens, which is the feature of the results most directly at odds with a reading in which the index simply registers the run-up to an imminent seizure.

Figure D.20: Out-of-Sample ROC Curves, Firm-Level Data versus IPI, Shorter Forecast Horizons



(a) One-Year-Ahead Prediction ($h = 1$)

(b) Two-Year-Ahead Prediction ($h = 2$)

Notes: As in Figure 3, for the two shortest forecast horizons. Green: firm-level data only. Red: Infiltration Proximity Index only.

D.8. Interpretation and Open Robustness Checks

The sampling frame The ownership network is built outward from judicially confirmed confiscation cases, so both classes are drawn from the same ecosystem: confiscated and non-confiscated firms in the sample tend to share geography, sector and relational environment. This cuts two ways. It makes the classification problem harder than the

Table D.25: Model Performance Comparison: Firm-level Data vs IPI Only, $h = 1$

<i>Panel A: Firm-level Data</i>				
	Logistic Regression	Random Forest	LightGBM	XGBoost
Accuracy	0.631	0.767	0.767	0.750
Precision	0.397	0.634	0.660	0.581
Recall	0.611	0.394	0.350	0.381
AUC	0.701	0.708	0.737	0.730
<i>Panel B: IPI Only</i>				
	Logistic Regression	Random Forest	LightGBM	XGBoost
Accuracy	0.806	0.795	0.797	0.777
Precision	0.611	0.601	0.588	0.567
Recall	0.847	0.803	0.919	0.853
AUC	0.895	0.866	0.903	0.884

Notes: Same sample, models, hyper-parameter search (RandomizedSearchCV, TimeSeriesSplit with 5 folds, ROC-AUC) and chronological train/test split as the main forecasting table. Panel A uses every firm-level variable at t and its lags $t-1, \dots, t-5$. Panel B uses only the Infiltration Proximity Index at t and its lags $t-1, \dots, t-4$, computed on the embeddings of the full-data TGN, whose node features include a time-varying confiscation indicator switched on from the confiscation year onward, with a time-respecting reference set (a neighbour counts as confiscated only from its confiscation year). Decisions at the models' 0.5 threshold. Bold values indicate the better information set for a given model and metric.

Table D.26: Model Performance Comparison: Firm-level Data vs IPI Only, $h = 2$

<i>Panel A: Firm-level Data</i>				
	Logistic Regression	Random Forest	LightGBM	XGBoost
Accuracy	0.659	0.770	0.756	0.747
Precision	0.411	0.611	0.584	0.553
Recall	0.601	0.408	0.344	0.332
AUC	0.718	0.662	0.720	0.722
<i>Panel B: IPI Only</i>				
	Logistic Regression	Random Forest	LightGBM	XGBoost
Accuracy	0.785	0.771	0.763	0.766
Precision	0.571	0.551	0.537	0.541
Recall	0.831	0.831	0.913	0.886
AUC	0.878	0.870	0.888	0.874

Notes: Same sample, models, hyper-parameter search (RandomizedSearchCV, TimeSeriesSplit with 5 folds, ROC-AUC) and chronological train/test split as the main forecasting table. Panel A uses every firm-level variable at t and its lags $t-1, \dots, t-5$. Panel B uses only the Infiltration Proximity Index at t and its lags $t-1, \dots, t-4$, computed on the embeddings of the full-data TGN, whose node features include a time-varying confiscation indicator switched on from the confiscation year onward, with a time-respecting reference set (a neighbour counts as confiscated only from its confiscation year). Decisions at the models' 0.5 threshold. Bold values indicate the better information set for a given model and metric.

Table D.27: Model Performance Comparison: Firm-level Data vs IPI Only, $h = 3$

<i>Panel A: Firm-level Data</i>				
	Logistic Regression	Random Forest	LightGBM	XGBoost
Accuracy	0.639	0.787	0.767	0.763
Precision	0.364	0.681	0.551	0.536
Recall	0.617	0.260	0.312	0.315
AUC	0.688	0.633	0.686	0.680
<i>Panel B: IPI Only</i>				
	Logistic Regression	Random Forest	LightGBM	XGBoost
Accuracy	0.777	0.745	0.737	0.737
Precision	0.532	0.491	0.483	0.483
Recall	0.820	0.804	0.878	0.907
AUC	0.881	0.849	0.881	0.876

Notes: Same sample, models, hyper-parameter search (RandomizedSearchCV, TimeSeriesSplit with 5 folds, ROC-AUC) and chronological train/test split as the main forecasting table. Panel A uses every firm-level variable at t and its lags $t-1, \dots, t-5$. Panel B uses only the Infiltration Proximity Index at t and its lags $t-1, \dots, t-4$, computed on the embeddings of the full-data TGN, whose node features include a time-varying confiscation indicator switched on from the confiscation year onward, with a time-respecting reference set (a neighbour counts as confiscated only from its confiscation year). Decisions at the models' 0.5 threshold. Bold values indicate the better information set for a given model and metric.

Table D.28: Model Performance Comparison: Firm-level Data vs IPI Only, $h = 4$

<i>Panel A: Firm-level Data</i>				
	Logistic Regression	Random Forest	LightGBM	XGBoost
Accuracy	0.729	0.793	0.804	0.812
Precision	0.436	0.613	0.695	0.739
Recall	0.549	0.304	0.287	0.297
AUC	0.746	0.651	0.723	0.711
<i>Panel B: IPI Only</i>				
	Logistic Regression	Random Forest	LightGBM	XGBoost
Accuracy	0.764	0.743	0.752	0.734
Precision	0.496	0.469	0.483	0.463
Recall	0.773	0.755	0.874	0.864
AUC	0.860	0.831	0.878	0.863

Notes: Same sample, models, hyper-parameter search (RandomizedSearchCV, TimeSeriesSplit with 5 folds, ROC-AUC) and chronological train/test split as the main forecasting table. Panel A uses every firm-level variable at t and its lags $t-1, \dots, t-5$. Panel B uses only the Infiltration Proximity Index at t and its lags $t-1, \dots, t-4$, computed on the embeddings of the full-data TGN, whose node features include a time-varying confiscation indicator switched on from the confiscation year onward, with a time-respecting reference set (a neighbour counts as confiscated only from its confiscation year). Decisions at the models' 0.5 threshold. Bold values indicate the better information set for a given model and metric.

one solved in studies that draw firms from the general population, which is the natural explanation for why our firm-only models underperform [Cariello et al. \(2024\)](#): within this ecosystem, standard financial indicators no longer separate the groups because the comparison firms are not average Italian firms. But it also means the reported figures are not an estimate of screening performance on the universe of Italian firms. What they establish is discrimination within a confiscation-centered network. Assessing operational screening performance would require the full population of ownership relations, to which we do not currently have access.

Full-data versus first-year index As noted in Section [D.3](#), the index is computed on the full-data model, whose event stream includes yearly financial self-loops. Repeating the exercise with the index computed on the first-year representation, and entering the two as separate information sets, would distinguish three information sources: firm variables only; an index summarizing both ownership and attribute histories; and an index driven exclusively by ownership events. Only the last would license the stronger claim that ownership dynamics alone forecast confiscation.

Further checks Three further robustness exercises would sharpen the reading of these results and are not implemented here: estimating the models without SMOTE, using class weights instead, to confirm that the rebalancing step is not doing the work; varying the number of lags of the index to establish how much of the signal comes from its level as opposed to its recent trajectory; and varying the matching rule used to assign synthetic reference years, since that rule determines the calendar-time alignment of the two classes. Because the precision–recall comparison is taken at a single decision threshold, reporting precision-recall curves alongside ROC curves, and the dispersion of metrics across random seeds and resamples, would also be informative given the class imbalance and the moderate sample size at longer horizons.