

# GreeSE Papers

## Hellenic Observatory Discussion Papers on Greece and Southeast Europe



**Paper No. 219**

### **Recurrence as a Governance Signal: Diagnostic Network Metrics for Public Procurement Oversight in Greece**

**Ioannis G. Fountoukidis, Eleni L. Dafli,  
Ioannis E. Antoniou and Nikos C. Varsakelis**

# **Recurrence as a Governance Signal: Diagnostic Network Metrics for Public Procurement Oversight in Greece**

**Ioannis G. Fountoukidis, Eleni L. Dafli, Ioannis E.  
Antoniou, and Nikos C. Varsakelis**

Papers by external academics who are not based at the LSE do not represent research carried out at the School. Each paper reflects the views and research findings of its respective author(s) and does not necessarily represent those of the Hellenic Observatory or the London School of Economics and Political Science.

GreeSE Paper No. 219

Hellenic Observatory Papers on Greece and Southeast Europe

# Recurrence as a Governance Signal: Diagnostic Network Metrics for Public Procurement Oversight in Greece

Ioannis G. Fountoukidis<sup>1</sup>, Eleni L. Dafli<sup>2</sup>, Ioannis E. Antoniou<sup>3</sup>, and Nikos C. Varsakelis<sup>4</sup>

## Abstract

This study examines what recurring buyer–supplier relationships in Greek public procurement reveal about routinised contracting behaviour. Drawing on contract award data from Greece's Central Electronic Registry of Public Contracts (KIMDIS, 2018–2025), it models procurement as a temporally evolving bipartite network across twelve domains, generating seven year-pair predictions. Three interpretable network signals — Historical Frequency, Preferential Attachment, and an adapted Adamic–Adar index — capture routinised continuity, structural concentration, and context-bound repetition. Predictability varies systematically across domains (AUC 0.80–0.96), with feature importance shifting from history-driven to structurally diverse recurrence over time. Authority-level analysis reveals extreme within-domain heterogeneity (vendor diversity 1–300+, HHI 0.008–1.0), demonstrating that uniform oversight thresholds are structurally inappropriate. The framework suggests differentiated governance responses — contestability reviews, dependency audits, and specification reform — and can be integrated into Greece's existing digital procurement infrastructure. Results are robust across three negative sampling specifications.

**Keywords:** public procurement, organisational routines, governance diagnostics, contestability, Greece.

## Data and Code Availability

All data and empirical results used in this study are publicly available via Zenodo (DOI: 10.5281/zenodo.20175804). The dataset includes the processed feature matrix and model outputs required to reproduce the analysis.

The replication code is available on GitHub at: <https://github.com/gallos3/procurement-network-recurrence>.

Due to file size constraints, the dataset is not included in the repository and must be downloaded separately from Zenodo.

---

<sup>1</sup> Corresponding author, Department of Economics Aristotle University of Thessaloniki, P.O.Box 184, 54124, Thessaloniki, Greece, E-Mail: [ifountouk@econ.auth.gr](mailto:ifountouk@econ.auth.gr)

<sup>2</sup> School of Medicine, Faculty of Health Sciences, Aristotle University of Thessaloniki, PO Box 376, 54124 Thessaloniki, Greece, E-Mail: [elendaf@yahoo.com](mailto:elendaf@yahoo.com)

<sup>3</sup> School of Mathematics, Faculty of Sciences, Complex Systems Analysis Laboratory (COSAL), Aristotle University of Thessaloniki, 541 24 Thessaloniki, Greece. E-Mail: [iantonio@math.auth.gr](mailto:iantonio@math.auth.gr)

<sup>4</sup> Department of Economics Aristotle University of Thessaloniki, P.O.Box 184, 54124, Thessaloniki, Greece E-Mail: [barsak@econ.auth.gr](mailto:barsak@econ.auth.gr)

## 1. Introduction

Public sector managers responsible for procurement face a persistent governance challenge: how to distinguish between beneficial continuity in buyer–supplier relationships and patterns of routinised contracting that may undermine competition, fairness, and accountability. Repeated awards to the same suppliers are a common feature of public procurement systems across jurisdictions (OECD, 2016; Thai, 2009). Such recurrence may reflect accumulated trust, specialisation, and administrative efficiency, but it may equally signal institutional inertia, dependency on incumbent suppliers, or diminished contestability. The managerial question is not whether recurrence exists — it clearly does — but what it reveals about institutional behaviour, and how public managers and oversight bodies can interpret and respond to it.

This question acquires particular salience in the Greek context. First, Greece has undertaken extensive procurement reforms over the past decade, establishing a modernised legal framework through Law 4412/2016 and creating the Hellenic Single Public Procurement Authority (EADHSY) for centralised oversight (Kontogeorga and Angelaras, 2023). Second, Greece's Central Electronic Registry of Public Contracts (KIMDIS) provides a uniquely comprehensive data infrastructure that records all public contracts regardless of value, including below-threshold awards where contracting authorities exercise greater discretion and where relational patterns are least visible to external oversight. Third, recent network-based research has revealed significant structural concentration in certain Greek procurement categories alongside more balanced competitive environments in others (Pliatsidis, 2024), pointing to substantial heterogeneity whose institutional mechanisms remain unexplored. Earlier work has also linked transparency conditions to procurement outcomes in Greece (Giotopoulos et al., 2015).

Research on recurrence in public procurement has developed largely outside the public management literature. Empirical studies have examined repeat contracting primarily through the lens of bilateral relationships and economic outcomes: some document efficiency gains from relational stability (Bandiera et al., 2009; Coviello and Mariniello, 2014; Yakovlev et al., 2015), while others highlight risks of favouritism and reduced competition (Palguta and Pertold, 2017; Baltrunaite et al., 2021), with recurring misconduct also documented as a systemic concern (Kistler et al., 2024). Studies on incumbent advantage show that contract renewal systematically favours existing suppliers (Albalade et al., 2020; Camboni and Valbonesi, 2020), and transparency and prior experience emerge as key determinants of incumbency persistence (Plaček et al., 2019). Analytical frameworks have been developed to measure corruption risk using procurement indicators (Fazekas and Tóth, 2016), and machine learning approaches have been applied for early detection of malfeasance (Gallego et al., 2021; Popa, 2019), while network analysis has examined competitive conditions (Fountoukidis et al., 2023; 2025) and corruption risk (Wachs et al., 2019).

A growing body of country-level procurement research has demonstrated the value of comprehensive national registries for uncovering governance dynamics that cross-

country studies obscure. Sturm et al. (2025) found that firms occupying influential network positions in Portuguese procurement consistently earn more per bid. Herrera et al. (2020) mapped community structures among suppliers in Chile's "Mercado Público". Fontana and d'Agostino (2025) detected manipulation of contract values below screening thresholds in Italian municipal procurement. Plaček et al. (2019) examined factors driving repeated supplier selection in the Czech Republic, and Kuo and Filz (2024) mapped collaborative communities in Finnish architectural services procurement. Yet none of these studies applies a diagnostic recurrence framework, one that uses the predictability of buyer–supplier relationships as a governance signal, decomposed into interpretable institutional mechanisms.

This body of work has three significant limitations from a public management perspective. First, it treats recurrence primarily as a bilateral phenomenon rather than as a system-level pattern reflecting broader institutional logics. Second, predictive and network-based approaches have been framed as tools for technical optimisation or corruption detection, with limited attention to their potential as diagnostic instruments for managerial oversight. Third, the connections between recurrence patterns and core public management concerns (accountability, routinised contracting behaviour, public value) have not been systematically articulated at the national level.

This paper addresses the following research question: What do recurring buyer–supplier relationships in Greek public procurement reveal about routinised contracting behaviour, and how can diagnostic analytics make these patterns interpretable for public managers?

To address this question, we reconceptualise recurrence as a governance signal, an observable pattern that reflects underlying institutional dynamics and can inform managerial judgment without prescribing procurement outcomes. Rather than treating prediction as a decision-automation tool, we use it diagnostically: the extent to which buyer–supplier relationships are predictable reveals how strongly institutional routines, supplier concentration, and context-specific factors shape procurement practice.

Empirically, the study draws on KIMDIS contract award data covering 2018–2025, modelling buyer–supplier relationships as a temporally evolving bipartite network across twelve procurement domains. Three interpretable network-based signals (Historical Frequency (HF), Preferential Attachment (PA), and an adapted Adamic–Adar (AA) measure) capture distinct dimensions of recurrence: routinised continuity, structural concentration, and context-bound repetition. The analysis examines recurrence at both the domain level and the contracting authority level, revealing substantial heterogeneity in institutional behaviour.

The paper makes three contributions. First, it provides one of the first applications of a temporal diagnostic recurrence framework to a national procurement registry covering the full contract value spectrum, demonstrating how different forms of recurrence correspond to distinct institutional logics that public managers can interpret and act

upon. Second, it advances the debate on data-driven public management by demonstrating how predictive analytics can function as reflective governance instruments rather than prescriptive mechanisms (Wirtz et al., 2019; Bannister and Connolly, 2020). Third, it establishes a multi-level analytical link between pairwise recurrence signals and authority-level market structure, including HHI, within a single national system, enabling comparative oversight that accounts for the extreme heterogeneity characteristic of Greek public procurement.

We interpret recurrence as a governance signal generated by the aggregation of authority-level routines and market structures. The framework does not assume a single, unified “memory” of the state. Rather, it makes visible how routinised continuity, hub-driven concentration, and context-bound matching vary across domains and authorities.

The remainder of the paper is organised as follows. Section 2 develops the conceptual framework. Section 3 describes the data, institutional context, and methodology. Section 4 presents the results, and Section 5 discusses the theoretical and managerial implications.

## **2. Conceptual Framework: Recurrence as a Governance Signal**

Recurring authority–supplier relationships are a normal feature of public procurement. Their governance meaning, however, is not fixed: the same observable repetition can reflect efficient learning and risk reduction, structural market constraints, or diminished contestability. This section links three streams of scholarship to clarify how recurrence is produced, why it matters, and how analytics can make it interpretable for oversight. Institutional theory explains how routines and path dependence generate recurrence (2.1). Public value theory provides the normative lens for judging when recurrence is benign or problematic (2.2). The literature on data-driven public management motivates a diagnostic use of prediction that supports reflection rather than decision automation (2.3).

### **2.1 Institutional Routines, Path Dependence, and the Production of Recurrence**

Procurement decisions unfold inside organisations that learn by doing. Formal rules shape what contracting authorities may do, but the practical conduct of procurement (how specifications are drafted, which suppliers are approached, which risks are treated as salient) depends on accumulated experience and relational familiarity. Over repeated contracting cycles, these practices sediment into institutional memory: embedded knowledge, relational configurations, and procedural habits that condition how authorities engage with markets (Levitt and March, 1988; March and Olsen, 1989).

In this paper, “institutional memory” is used in its organisational sense: it refers to routines, procedural habits, and relational familiarity that develop within contracting authorities over repeated procurement cycles (Levitt and March, 1988; March and Olsen, 1989). Our empirical unit of analysis is the authority–supplier pair, and the mechanisms we infer (HF/PA/AA) are therefore best interpreted as *authority-*

*level* organisational routines and market-structure effects. When we describe “system-level” recurrence patterns, we use the term to denote the aggregate distribution of authority-level routines across the procurement system rather than a single, unified memory of the state.

Path dependence explains why such configurations persist (Pierson, 2000). When an authority and supplier complete a contract successfully, several reinforcing dynamics follow. Authorities gain information that reduces perceived risk in re-engagement. Administrative processes adapt to the incumbent’s operational profile. Transaction costs fall because the parties know each other’s routines and cognitive shortcuts favour familiar options. In technically complex or regulated domains, or where supplier pools are thin, these dynamics can represent an efficient institutional adaptation rather than a governance failure.

Yet recurrence can also entrench incumbency. Prior work shows that contract renewal and repeated selection may systematically advantage incumbents, consolidate market positions, and reduce competitive opportunity even under formally open procedures. The same repetition can therefore signal different underlying mechanisms: learning and risk management, structural constraints, or institutional rigidity. The governance problem is not to “eliminate recurrence,” but to interpret it in a way that distinguishes these mechanisms.

## **2.2 Public Value as the Normative Lens for Recurrence**

Institutional theory clarifies how recurrence arises. Public value theory clarifies why it demands governance attention. Public organisations are accountable not only for economic efficiency but also for equitable access to public markets, transparency in allocating public resources, competitive fairness, and public trust. Procurement practices should therefore be evaluated against a portfolio of public value criteria rather than cost alone, a point that complements the efficiency focus of classic procurement theory (Laffont and Tirole, 1993).

Recurrence matters because it can move these criteria in opposite directions. Stable authority–supplier ties may increase public value when they preserve specialised knowledge, sustain reliable delivery, or reduce administrative burden in repeated purchasing. The same stability may erode public value when it forecloses access for capable competitors, creates dependency that weakens performance discipline, or renders award patterns opaque to oversight.

The normative assessment is contextual. In a thin market with few certified suppliers, high recurrence may be structurally inevitable and compatible with public value. In a market with many alternatives, similar recurrence may indicate diminished contestability. This perspective also differentiates the present study from adjacent research traditions. Corruption-focused work often treats repetition as a red flag requiring investigation, while economic work treats it as a bilateral outcome whose welfare implications depend on price and efficiency effects. A public management

perspective treats recurrence as a system property that can be governance-relevant even when it is neither corrupt nor price-distorting (Peters, Pierre and Randma-Liiv, 2011).

### **2.3 From Predictive Analytics to Diagnostic Governance**

Governments increasingly use data analytics in procurement oversight. The most common applications either train algorithms to flag potentially corrupt transactions or use network analysis to visualise market structure. These tools can be valuable, but they share a common logic: they classify. A transaction is flagged or not, a market is concentrated or not. The underlying question is always whether something looks wrong.

This approach, however, carries an underappreciated risk. When predictive models learn from historical procurement data, they inevitably absorb the relational patterns that data contains, including any entrenched advantages or exclusionary tendencies. If such models are then used to recommend suppliers, generate shortlists, or trigger automatic reviews, they risk reinforcing the very patterns they were meant to scrutinise. In practice, this means that old habits and established power structures get preserved, not because anyone chose to keep them, but because the algorithm treats them as the baseline for what is "normal." Research on algorithmic governance in the public sector has repeatedly highlighted this danger (Mittelstadt et al., 2016; Ziewitz, 2016; Veale and Binns, 2017).

This study adopts a diagnostic alternative. We use prediction to measure how strongly procurement outcomes are structured by prior relational history, market centrality, and contextual matching. In this framing, the degree of predictability and, crucially, the sources of predictability become governance information. When prediction is driven primarily by Historical Frequency, recurrence reflects routinised continuity: relationships reproduce themselves through institutional memory and habit. When prediction is driven by Preferential Attachment, recurrence reflects hub-dependent concentration: structurally prominent suppliers attract awards because of centrality and scale. When prediction is driven by an adapted Adamic–Adar signal, recurrence reflects context-bound repetition: buyers and suppliers repeatedly co-occur in specialised procurement niches shaped by specification and context.

The point of this approach is not to tell oversight bodies what to do, but to help them ask the right questions. In a given domain, is recurrence driven mainly by habit? By a few dominant suppliers? Or by the technical specificity of what is being purchased? Each answer points to a different response: reviewing whether competition is being tested, assessing whether dependency on key suppliers has become excessive, or examining whether procurement specifications are unnecessarily narrow. The framework provides a structured way to diagnose, not a one-size-fits-all threshold for intervention.

The Greek institutional ecosystem makes this approach practically relevant. KIMDIS records the full spectrum of awards, ESIDIS captures procedural execution, and EADHSY holds the oversight mandate. The missing element is an analytical layer that translates

registry data into interpretable governance intelligence. The framework developed here shows how to construct that layer using transparent network-based signals. These institutional considerations are reflected in the methodological choices described in Section 3, in particular the decision to use interpretable network features and to present results as governance signals for practitioner review rather than as decision-automation outputs.

### **3. Data and Methodology**

#### **3.1. Data source and graph construction**

The empirical analysis draws on contract award data from the Central Electronic Registry of Public Contracts (KIMDIS), the official Greek platform for the mandatory registration of all public procurement contracts (Law 4412/2016, Art. 38). KIMDIS records encompass the full range of public contracting activity in Greece, including contracts both above and below EU publication thresholds. This distinguishes the dataset from supra-national sources such as Tenders Electronic Daily (TED), which captures only contracts exceeding EU-mandated value thresholds and therefore omits the substantial volume of lower-value procurement where contracting authorities exercise greater discretion.

The dataset spans the years 2018 to 2025 and includes information on awarded contracts, contracting authorities, suppliers, contract values, Common Procurement Vocabulary (CPV) codes, and award dates. Contract records are structured around three core entities: contracting authorities (public bodies issuing procurement), economic operators (firms or consortia receiving awards), and contract objects (classified by CPV code).

Procurement relationships are represented as a bipartite graph connecting contracting authorities to suppliers through contract award edges. Each edge carries temporal and economic annotations: the year of award, the contract value, and the CPV code identifying the procurement domain. Contract value is used in the authority-level market-structure indicators (HHI, top-supplier share) after data-quality screening, while the link-prediction features are deliberately count-based to preserve robustness and interpretability.

The resulting temporal graph encompasses 3,265 contracting authorities, 224,539 companies, and 2,195,473 contract awards registered in KIMDIS over the period 2018–2025, enabling longitudinal analysis of how buyer–supplier relationships form, persist, and dissolve across the full spectrum of Greek public procurement.

#### **3.2 Domain Selection**

To balance analytical depth with sectoral breadth, we select twelve five-digit CPV categories for detailed analysis. The selection procedure ensures that the chosen domains are both empirically tractable (with sufficient relational density to support network-based signals) and substantively diverse, spanning different market structures,

regulatory environments, and contracting patterns characteristic of the Greek public sector.

The selection follows a structured protocol. First, we retain only CPV categories that appear consistently across all eight observation years (2018–2025), excluding domains with discontinuous reporting that would compromise temporal analysis. Second, for each eligible category, we compute the total number of contract awards and the total awarded value across the observation period. Third, we construct a composite ranking score that prioritises interaction density while retaining economic relevance:

$$S_i = 0.7 \cdot \text{Rank}(N_i) + 0.3 \cdot \text{Rank}(\log(1+V_i))$$

where  $N_i$  is the total number of contracts and  $V_i$  the total awarded value for CPV category  $i$  over the observation period. We use  $\log(1+V_i)$  to reduce the influence of a small number of very high-value awards. The heavier weight on contract counts reflects the requirements of network-based inference: link-prediction signals depend primarily on repeated interactions, so domains must be sufficiently dense to support stable feature construction. The value term acts as a secondary guardrail against selecting domains that are high-frequency but economically marginal. Sensitivity checks under alternative weights (0.5/0.5 and 0.8/0.2) yield a highly similar top-12 set, indicating that the domain selection is not driven by the exact weighting choice.

Table 1 presents the selected categories, their CPV codes, and summary structural characteristics.

*Table 1: Selected Procurement Domains*

| cpv   | Description                                   | n_contracts | n_authorities | n_suppliers |
|-------|---|-------------|---------------|-------------|
| 30192 | Office machinery, equipment and supplies      | 27724       | 1638          | 4377        |
| 33111 | Radiography equipment                         | 12534       | 152           | 1881        |
| 33140 | Medical consumables                           | 98808       | 492           | 5798        |
| 33141 | Minor medical devices and products            | 62643       | 472           | 4753        |
| 33696 | Reagents and contrast media                   | 68750       | 353           | 2377        |
| 44111 | Building materials                            | 15316       | 1112          | 4446        |
| 45233 | Road construction and repair works            | 23427       | 710           | 5238        |
| 50000 | Repair and maintenance services               | 23404       | 1104          | 7912        |
| 50112 | Motor vehicle maintenance and repair services | 32896       | 689           | 4447        |
| 79713 | Guard and security services                   | 11764       | 877           | 847         |
| 90910 | Cleaning services                             | 16105       | 1096          | 2896        |
| 90911 | Building cleaning services                    | 12241       | 786           | 1474        |

The selected domains span repair and maintenance services (50112, 50000), medical and pharmaceutical supplies (33696, 33140, 33141, 33111), construction materials and road works (44111, 45233), office and computing supplies (30192), cleaning services (90910, 90911), and security services (79713). This composition captures a broad cross-

section of the Greek procurement landscape, encompassing both highly specialised markets with limited supplier pools and more open, commoditised markets with diverse competitive conditions.

### 3.3 Feature Engineering: Network-Based Recurrence Signals

For each CPV domain and base year  $t$ , we construct a set of candidate authority–supplier pairs and compute three interpretable, network-based recurrence signals using only award history observed up to and including  $t$ . The signals draw on standard link-prediction ideas, but are chosen here for interpretability: each corresponds to a distinct mechanism through which repeated contracting can arise in procurement networks.

**Historical Frequency (HF).** HF captures direct relational persistence by counting the number of distinct years in which an authority–supplier pair has an observed award relationship in the given CPV domain up to the base year:

$$HF(a, c) = \sum_{\tau=2018}^t 1\{\text{contract}(a, c, \tau)\}$$

As a summary of prior pair-specific activity, HF is expected to be strongly predictive of future recurrence even under strict temporal separation. In this study we therefore treat HF primarily as a descriptive, diagnostic indicator of routinised continuity rather than as a causal explanation.

**Preferential Attachment (PA).** PA captures hub-dependent concentration by measuring the product of the CPV-scoped degrees of the authority and the supplier (Barabási and Albert, 1999):

$$PA(a, c) = \text{deg}(a) \times \text{deg}(c)$$

where  $\text{deg}(\cdot)$  denotes the number of unique contracting partners within the CPV domain observed up to year  $t$ . Higher PA values indicate that recurrence is associated with structurally central actors (large authorities and/or dominant suppliers), consistent with concentration dynamics driven by network centrality rather than pair-specific history.

**Adapted Adamic–Adar (AA).** To capture context-bound repetition, we use an Adamic–Adar–inspired rarity-weighted co-occurrence score defined over award events (record-level awards) in the bipartite procurement graph. Let  $C(a, c)$  denote the set of past award events in the CPV domain (up to year  $t$ ) that link authority  $a$  and supplier  $c$ . We define

$$AA(a, c) = \sum_{z \in C(a, c)} \frac{1}{\log |N(z)|}$$

where each  $z$  is an award-event node (a specific award record) connecting  $a$  and  $c$ , and  $|N(z)|$  is the number of distinct entities linked to that award record in the graph (e.g., the participating authority and supplier and any additional contextual nodes recorded with the award). This adaptation differs from the standard unipartite Adamic–Adar index, which sums over shared neighbours. Here, the “intermediaries” are award events, and rarity is operationalised as low event connectivity: events that are connected to fewer entities contribute more weight. Intuitively, repeated co-occurrence through such “niche” events receives greater weight than repetition occurring through ubiquitous, widely connected events, making AA a proxy for context-bound matching rather than general relational frequency.

Contract value is excluded from the link-prediction features for two reasons. First, value data in KIMDIS — as in most national registries — contains missing entries, reporting errors, and multi-CPV awards whose total value cannot be reliably attributed to a single domain. Count-based signals (HF, PA, AA) are therefore more robust for prediction purposes. Second, and more fundamentally, the three features are designed to detect the *structural pattern* of co-occurrence — whether and how often a pair interacts — which is conceptually distinct from the economic scale of individual transactions. Value enters the analysis where it is both interpretable and attributable: at the authority level, after outlier screening, through HHI, top-supplier share, and vendor diversity. This design separates the structural recurrence signal (pair-level, count-based) from the market-structure signal (authority-level, value-based), keeping each layer internally consistent.

**Authority-level contextual indicators.** To situate pairwise recurrence signals within broader market structure, we compute annual, CPV-scoped summaries for each contracting authority: (i) *vendor diversity*, the count of distinct suppliers receiving awards from the authority within the domain-year, (ii) *top-supplier share*, the ratio of the highest single-supplier contract value (net of VAT) to the authority’s total expenditure within the domain-year, and (iii) the Herfindahl–Hirschman Index (HHI), defined as

$$HHI_a = \sum_{c=1}^{c_a} \left( \frac{v_{ac}}{V_a} \right)^2$$

where  $v_{ac}$  is total contract value awarded by authority  $a$  to supplier  $c$  in the domain-year,  $V_a$  is total domain-year expenditure by authority  $a$ , and  $c_a$  is the number of distinct suppliers. Together these indicators capture complementary aspects of authority-level market structure (breadth of the supplier base and concentration of spending) and enable within-domain comparisons that reveal heterogeneity masked by aggregate domain statistics.

### 3.4 Model Specification and Evaluation

We employ XGBoost (Chen and Guestrin, 2016) as the classification algorithm. Three properties motivate this choice for a diagnostic study. First, gradient-boosted trees

accommodate the sparsity and class imbalance inherent in procurement networks without requiring extensive data preprocessing. Second, XGBoost's gain-based feature importance decomposition directly supports the comparative analysis of institutional mechanisms, quantifying how much each signal (HF, PA, AA) contributes to predictive discrimination across domains and over time. Third, the tree-based architecture captures non-linear feature interactions while preserving the interpretability of individual input variables, occupying a middle ground between fully transparent linear models and opaque deep learning approaches.

The choice of XGBoost over simpler linear specifications reflects the diagnostic requirements of this study. Gradient-boosted trees capture non-linear feature interactions while preserving the ability to decompose each prediction into the individual contributions of HF, PA, and AA through gain-based importance.

**Temporal design.** For each base year  $t$ , features are computed using only award history observed up to and including  $t$ , while the outcome is link presence in  $t+1$ , enforcing strict temporal ordering and preventing information leakage. Performance is evaluated on held-out authority–supplier pairs within each  $t \rightarrow t+1$  dataset. Separate models are estimated for each CPV category, allowing domain-specific structural dynamics to emerge rather than being averaged across heterogeneous markets. Only authorities and CPV categories active in both  $t$  and  $t+1$  are retained, ensuring that recurrence is assessed within a stable institutional perimeter.

**Class balance.** Actual contract awards represent a small fraction of all possible authority–supplier combinations within a CPV domain. To manage this imbalance, we draw negative samples (authority–supplier pairs with no historical connection up to year  $t$ ) at three controlled ratios: 10%, 50%, and 100% of the positive class, with a ceiling of 100,000 negative samples per run. Evaluating across all three ratios distinguishes predictive patterns rooted in genuine institutional structure from artefacts of sampling design. In our setting, the 100,000 ceiling was not binding in any CPV×year-pair run. Effective negative ratios equal the nominal ratios throughout (Appendix, Table A3). Domain rankings by AUC and dominant feature profiles are stable across all three ratios (Appendix, Table A4).

**Evaluation metrics.** Model performance is assessed through precision, recall, F1-score, and AUC. Beyond standard interpretation, these metrics carry governance meaning: precision reflects how tightly institutional patterns constrain outcomes. Recall indicates how much recurrence escapes established structural configurations, and AUC captures the overall structural legibility of recurrence within a domain.

### 3.5 Safeguards and Institutional Context

The methodological design embeds explicit safeguards against the misuse of predictive outputs. Model results are conceived as inputs for managerial reflection and oversight prioritisation, not as automated decision criteria. Their intended users are oversight bodies such as EADHSY, audit institutions such as the Hellenic Court of Auditors, and

strategic procurement units within contracting authorities. The purpose is to identify domains and authorities where recurrence patterns warrant contextual investigation, not to prescribe procurement outcomes.

The Greek legal framework provides mechanisms (framework agreements (Art. 39, Law 4412/2016), Dynamic Purchasing Systems (Art. 40), and centralised purchasing bodies) that legitimately structure repeated contracting. Recurrence diagnostics must be interpreted with awareness of these mechanisms: elevated HF within a framework agreement may reflect institutional design rather than failure.

Diagnostic thresholds are calibrated at the domain level using authority-level indicators (vendor diversity, top-supplier share and HHI), ensuring that governance interpretations are anchored in market-specific conditions rather than arbitrary cutoffs.

## 4. Results: Predictive Patterns and Institutional Signals

### 4.1 Overview of Predictive Performance

Table 2 reports mean predictive performance across twelve CPV domains under the primary specification (negative sampling ratio = 1.0), averaged over seven year-pairs (2018→2019 through 2024→2025). Full year-pair results are reported in Appendix Table A3.

*Table 2: Predictive performance by CPV category and year-pair (negative ratio = 1.0)*

| CPV Code | Category                      | Predictive Performance |        |                |             |         | Rank     |         |
|----------|-------------------------------|------------------------|--------|----------------|-------------|---------|----------|---------|
|          |                               | Mean AUC               | SD AUC | Mean Precision | Mean Recall | Mean F1 | AUC Rank | F1 Rank |
| 30192    | Office Supplies               | 0.869                  | 0.008  | 0.505          | 0.204       | 0.288   | 7        | 11      |
| 33111    | Radiography Equipment         | 0.912                  | 0.053  | 0.584          | 0.449       | 0.505   | 4        | 3       |
| 33140    | Medical Consumables           | 0.935                  | 0.006  | 0.651          | 0.454       | 0.531   | 2        | 2       |
| 33141    | Minor Medical Devices         | 0.923                  | 0.006  | 0.614          | 0.384       | 0.464   | 3        | 4       |
| 33696    | Reagents & Lab Chemicals      | 0.939                  | 0.005  | 0.703          | 0.608       | 0.644   | 1        | 1       |
| 44111    | Building Materials            | 0.866                  | 0.015  | 0.520          | 0.267       | 0.351   | 8        | 7       |
| 45233    | Road Construction/Repair      | 0.865                  | 0.013  | 0.501          | 0.225       | 0.295   | 9        | 9       |
| 50000    | Repair & Maintenance Services | 0.874                  | 0.009  | 0.474          | 0.166       | 0.246   | 6        | 12      |
| 50112    | Motor Vehicle Maintenance     | 0.907                  | 0.021  | 0.600          | 0.352       | 0.439   | 5        | 5       |
| 79713    | Guard & Security Services     | 0.845                  | 0.018  | 0.481          | 0.252       | 0.333   | 12       | 8       |
| 90910    | Cleaning Services             | 0.859                  | 0.016  | 0.501          | 0.210       | 0.294   | 10       | 10      |
| 90911    | Building Cleaning Services    | 0.857                  | 0.018  | 0.502          | 0.279       | 0.333   | 11       | 6       |

Two key patterns emerge from the results. First, discrimination is consistently strong: across all domain–year combinations, AUC ranges from 0.80 to 0.96. Second, point prediction exhibits an asymmetry between precision and recall. Precision averages roughly 0.55 (about one in two predicted recurrences corresponds to an actual repeat tie), while recall is lower and more variable (mean  $\approx$  0.31 with range  $\approx$  0.09–0.76). In

other words, the model is good at telling which pairs are more likely to recur than others (high AUC), but many actual recurrences happen for reasons that the three network signals alone do not capture (lower recall).

Performance differs systematically across domains. Medical and pharmaceutical procurement displays the strongest and most stable predictability. CPV 33696 achieves the highest average AUC (0.94) and the most balanced precision–recall tradeoff, with F1 scores ranging from 0.57 to 0.76 across the seven year-pairs. CPV 33140 displays a similar pattern, with AUC values consistently between 0.93–0.94 and F1 between 0.47–0.64. CPV 33141 follows closely, with AUC around 0.91–0.93 and F1 between 0.35–0.65. These domains are characterised by certification requirements, regulatory constraints on supplier qualification, and centralised purchasing, conditions that structurally favour relational continuity.

CPV 33111 (medical imaging equipment) is a distinctive case among medical categories. It records the widest AUC range of any domain (0.80 in 2018→2019 rising to 0.96 in 2024→2025) and the single highest F1 score in the dataset (0.78 in 2019→2020), but also pronounced year-to-year volatility, consistent with a thin, specialised market where a small number of high-value contracts can shift metrics substantially.

Repair and maintenance domains occupy an intermediate position, with clear deterioration over time. CPV 50112 (vehicle repair) starts with strong performance (F1 = 0.65, AUC = 0.94 in 2018→2019) but deteriorates progressively, reaching F1 = 0.30 and AUC = 0.88 by 2024→2025. CPV 50000 (general repair) follows an even steeper decline, from F1 = 0.38 down to F1 = 0.19.

The weakest point-prediction performance characterises construction, supplies, and service domains. CPV 45233 (road construction) maintains moderate AUC (0.85–0.88) but very low recall, producing F1 scores that fall from 0.46 in 2018→2019 to 0.18 in 2024→2025. CPV 44111, 79713, 30192, and cleaning (90910/90911) exhibit similar profiles: AUC around 0.83–0.88 but F1 predominantly below 0.40.

This finding may appear to contradict the well-documented observation that construction and cleaning contracts are among the most prone to repeated awards to the same suppliers in Greece and elsewhere. The apparent tension, however, is informative rather than contradictory. The model's high AUC confirms that recurrence in these domains is structurally patterned: the network signals do separate recurring from non-recurring pairs. What the low recall reveals is that much of this recurrence is driven by factors that lie outside the three network features: specification design that effectively narrows the eligible supplier pool, geographic constraints, framework agreements, or informal practices that channel awards toward familiar contractors. These are precisely the mechanisms that audit reports and anti-fraud investigations typically identify and precisely the mechanisms that registry data alone cannot encode. The diagnostic value of the finding lies in this gap: domains where recurrence is high but

poorly explained by observable network structure are domains where non-structural drivers, whether legitimate or problematic, deserve the closest oversight attention.

Temporal dynamics reinforce this heterogeneity. Across most domains, recall declines over successive year-pairs, while AUC remains relatively stable (and in some domains improves). This divergence suggests that the relative ordering of pairs remains informative (discrimination persists), but the boundary separating recurring from non-recurring ties becomes less distinct over time. The seven-year window also reveals that deterioration is not uniformly monotonic: several domains show partial recovery after 2022, consistent with external shocks (notably COVID-era disruption) overlaying longer-run structural change.

The seven-year window allows a more nuanced reading of these trajectories. In CPV 90911, recall dips to 0.18 in 2024→2025 but had partially recovered to 0.28 in 2023→2024 after a trough of 0.20 in 2022→2023. In CPV 44111, F1 rises from 0.31 in 2020→2021 to 0.45 in 2022→2023 before declining again. These non-monotonic patterns suggest that COVID-related supply chain disruptions (visible in the 2020→2021 and 2021→2022 periods) temporarily reshuffled buyer–supplier relationships, and whether post-2022 recovery represents a return to pre-pandemic relational patterns or the emergence of new institutional configurations remains an open empirical question.

Robustness checks confirm that the core diagnostic picture does not depend on the negative sampling specification. Domain rankings by AUC remain stable across ratios, and qualitative feature-importance profiles (which mechanisms dominate in which domains) are preserved (Appendix, Table A1). As expected, precision/recall shift mechanically with class balance, AUC—being threshold-independent—provides the most comparable performance signal across sampling regimes (Appendix, Table A4).

#### **4.2 Feature Contributions: Sources of Recurrence**

To interpret recurrence as a governance signal, we decompose each model using gain-based feature importance across the three interpretable signals: Historical Frequency (HF), Preferential Attachment (PA), and the adapted Adamic–Adar index (AA). Gain captures the marginal contribution of each signal to model performance, allowing comparisons across domains and over time. Because gain-based importance can be sensitive to model complexity, we conducted a targeted hyperparameter sensitivity check (varying tree depth and learning rate on representative CPV domains and early/late year-pairs). Qualitative HF/PA/AA mechanism profiles are largely stable, and the few dominant-feature switches occur only when HF and AA are near-tied (Appendix, Table A2).

A consistent temporal pattern emerges: HF dominates early year-pairs, while PA and AA contributions rise as the observation window expands and recurrence becomes explainable by more than prior repetition. In CPV 33696 (laboratory reagents), HF accounts for 76.6% of gain in 2018→2019 versus only 8.5% for PA and 14.9% for AA. By 2024→2025, HF falls to 46.4% while AA rises to 38.0% and PA to 15.6%. This

redistribution, from simple relational persistence toward a mix of structural and contextual signals, is not a methodological curiosity. It changes how recurrence should be interpreted.

Domain profiles differ in diagnostically meaningful ways.

HF-dominant predictability indicates routinised continuity: recurrence reflects organisational routines (institutional memory in the organisational sense) and established relational configurations. This pattern is most visible in medical and pharmaceutical categories, where repeated contracting is compatible with regulated supplier qualification and centralised purchasing architectures. Yet even in CPV 33140 (medical consumables), the mechanism mix evolves: by 2024→2025, AA accounts for 45.9% of gain while HF falls to 34.9%, suggesting that context-bound matching increasingly supplements relational persistence.

PA-heavy predictability indicates hub-dependent concentration: recurrence follows the gravitational pull of structurally central suppliers, consistent with markets shaped by scale economies, dominant distributors, or entrenched intermediaries.

AA-heavy predictability indicates context-bound repetition: recurrence concentrates in specialised, relatively rare procurement contexts within a domain, suggesting that specification and niche matching, rather than simple relational inertia, structure repeat ties.

Construction and service domains illustrate why this decomposition matters. In CPV 45233 (road construction), the first year-pair is almost entirely HF-driven ( $hf\_gain = 0.946$ ), but by 2023→2024 the model assigns comparable weight to HF and AA (0.380 and 0.408 respectively), with PA also non-trivial (0.213). Similarly, CPV 90910 (cleaning) shifts from HF-dominance (0.944) in 2018→2019 to an AA-led regime by 2022→2023 ( $aa\_gain = 0.420$ ;  $pa\_gain = 0.194$ ;  $hf\_gain = 0.385$ ). Even when recall is low, these shifts imply that recurrence remains structured but conditional and context-specific. Low recall should not be read as "no institutional structure," but as a sign that a larger share of recurrences is driven by factors not encoded in the three signals.

CPV 33111 (medical imaging equipment) remains an instructive outlier, combining very high discrimination in later periods ( $AUC = 0.952$  in 2024→2025) with a mechanism mix in which AA remains prominent (0.376) alongside HF (0.425) and PA (0.200), consistent with a thin, specialised market where niche matching and structural concentration drive recurrence rather than simple repetition.

A simple comparison of average feature values confirms these patterns. Pairs that do recur tend to have much higher PA and AA scores than pairs that do not, meaning they involve better-connected suppliers and operate in more specialised procurement contexts. For example, in CPV 33140 (medical consumables, 2024→2025), recurring pairs have an average PA of 9,282 compared to just 1,054 for non-recurring pairs, and

an average AA of 15.93 versus 0.74. The gap in HF is smaller (3.25 versus 0.64), which is expected: HF counts prior years of interaction and is capped by the length of the observation window.

### 4.3 Authority-Level Market Structure and Recurrence Mechanisms

The pairwise recurrence signals examined in Sections 4.1 and 4.2 operate at the level of individual authority–supplier relationships. To connect these signals with broader market structure, we examine whether the relative importance of HF, PA, and AA as predictive mechanisms varies systematically with authority-level concentration indicators. This directly addresses the diagnostic claim that the two analytical layers—pairwise mechanism profiles and authority-level market structure—are complementary rather than parallel.

*Table 3. Spearman Rank Correlations: Feature Importance and Predictive Performance vs. Authority-Level Market Structure (CPV × Year-Pair Level, n = 84)*

|         | HHI       | Vendor Diversity | Top-Supplier Share |
|---------|-----------|------------------|--------------------|
| HF Gain | 0.200†    | -0.277*          | 0.209†             |
| PA Gain | -0.163    | 0.268*           | -0.180             |
| AA Gain | -0.178    | 0.196†           | -0.177             |
| AUC     | -0.710*** | 0.763***         | -0.722***          |

*Notes: Spearman  $\rho$  coefficients. Each observation is one CPV–year-pair (12 CPV domains × 7 year-pairs, one excluded due to single-class outcome; n = 84). Market structure variables are domain-year means computed from the authority-level heterogeneity dataset. HF = Historical Frequency gain; PA = Preferential Attachment gain; AA = Adamic–Adar gain; AUC = area under the ROC curve (negative sampling ratio = 1.0). † $p < 0.10$ ; \* $p < 0.05$ ; \*\*\* $p < 0.001$ .*

Across 84 CPV×year-pair observations, AUC correlates negatively with HHI (Spearman  $\rho = -0.71, p < 0.001$ ) and top-supplier share ( $\rho = -0.72, p < 0.001$ ), and positively with vendor diversity ( $\rho = +0.76, p < 0.001$ ). This implies that recurrence is most structurally legible—most predictable from observable network features—in more competitive, diverse settings. In highly concentrated markets, a larger share of recurrence appears to operate through channels not captured by the three network signals, consistent with the non-structural drivers discussed in Section 4.1.

Mechanism profiles shift with market structure in the expected direction. HF gain correlates negatively with vendor diversity ( $\rho = -0.28, p = 0.011$ ), indicating that history-driven continuity becomes more prominent where supplier pools are narrow. PA gain correlates positively with vendor diversity ( $\rho = +0.27, p = 0.014$ ), suggesting that hub-dependent dynamics become more salient as participation broadens. Grouping domains by concentration makes the contrast concrete: Very High concentration domains (mean HHI > 0.75) exhibit higher mean HF gain (0.528 vs 0.457) and lower mean PA gain (0.156 vs 0.198) than High (0.50–0.75) concentration domains.

Taken together, these results empirically connect the two analytical layers. Domain-level mechanism profiles are not arbitrary model artefacts: they align with independently observable authority-level market structure and can therefore be used jointly in a two-dimensional oversight framework.

## **5. Discussion**

### **5.1 Theoretical Contributions**

The analysis is among the first to apply a temporal, diagnostic recurrence framework to a national procurement registry that covers the full contract-value spectrum, including below-threshold awards. It combines temporal link prediction with an interpretable decomposition of recurrence mechanisms (HF/PA/AA). Prior work has often relied on supra-national databases limited to above-threshold contracts or on narrowly scoped national datasets. By encompassing awards both above and below EU thresholds, the KIMDIS-based analysis captures the complete contracting landscape of a single member state, including the lower-value transactions where procedural discretion is greatest and relational patterns are least visible.

Substantively, the results move the discussion beyond the binary question of whether recurrence exists in procurement (Plaček et al., 2019; Albalade et al., 2020; Camboni and Valbonesi, 2020) toward the more consequential question of how different forms of recurrence can be identified, measured, and governed. The finding that AUC ranges from 0.80 in fragmented construction markets to 0.96 in thin medical equipment markets (and that these rankings are invariant across three negative sampling specifications) demonstrates that recurrence patterns are not uniform across procurement domains. Their strength and form vary with market structure, regulatory environment, and the nature of the goods and services procured. This analytical reach positions the study within an emerging tradition of country-level procurement research that exploits comprehensive national registries to examine governance dynamics at a depth cross-country comparisons cannot achieve (Sturm et al., 2025; Herrera et al., 2020; Fontana and d'Agostino, 2025).

The second contribution is methodological. The study validates a diagnostic framework in which predictive analytics function as reflective governance instruments. The correspondence between each feature and a specific institutional mechanism gives oversight actors a vocabulary for interpreting what drives recurrence in a given domain, rather than merely knowing that recurrence exists.

The third contribution connects micro-level relational patterns to meso-level market structure through the authority-level heterogeneity analysis. By situating pairwise recurrence signals within the distribution of vendor diversity, top-supplier share, and HHI at the authority level, the study reveals heterogeneity that aggregate domain statistics obscure. The authority-level distributions provide concrete empirical grounding for public value claims about equitable market access and proportionate oversight (Moore, 1995; Bozeman, 2002; Bryson, Crosby and Bloomberg, 2014): they

demonstrate that uniform thresholds are structurally inappropriate for procurement systems characterised by such extreme within-domain variation.

## 5.2 Managerial Implications

The findings suggest a differentiated *diagnostic* framework for procurement oversight. The proposed mappings from diagnostic signals to governance responses are illustrative hypotheses intended to guide prioritisation and further investigation, rather than prescriptive rules.

Authorities exhibiting HF-dominant recurrence, high relational persistence coupled with limited vendor diversity and elevated HHI, display the characteristics of routinised contracting. The governance question is whether established relationships reflect justified efficiency or diminished contestability. EADHSY can leverage these signals to prioritise audit attention toward authorities whose recurrence profiles deviate from domain-specific norms. For example, an authority exhibiting HHI above 0.7 in a domain where comparable authorities routinely achieve HHI below 0.1 would be a plausible candidate for prioritised contestability review, not because high concentration is inherently problematic, but because it is anomalous relative to demonstrated market conditions.

PA-driven concentration, where awards cluster around structurally prominent suppliers, indicates dependency dynamics that may constrain market access. These patterns can inform dependency assessments and supplier-onboarding initiatives, with direct relevance to Greece's EU obligations regarding SME participation.

AA-driven contextual repetition, characterised by broader diversity but specification-driven recurrence, suggests that a reasonable starting point for review is specification practices rather than market structure: standardisation, functional specification, or knowledge-sharing among authorities facing similar procurement challenges. The temporal shift documented in Section 4.2, in which AA contributions increase progressively across the seven year-pairs, suggests that specification-oriented governance responses will grow in importance.

The authority-level heterogeneity documented in Section 4.3 sharpens these implications. The coexistence of authorities with vendor diversity of 1 and HHI of 1.0 alongside authorities with over 300 suppliers and HHI below 0.01, within the same CPV category and year, means that any governance response applied uniformly across a domain will be structurally miscalibrated. EADHSY and audit bodies can use the combination of domain-level mechanism profiles (HF/PA/AA) and authority-level market-structure indicators (vendor diversity, top-supplier share, HHI) to construct a two-dimensional oversight matrix: the first dimension identifies what drives recurrence in a domain. The second dimension identifies which authorities deviate from domain norms. This approach enables proportionate oversight that concentrates audit resources where they are most likely to reveal actionable governance information, for

instance, on authorities exhibiting high HHI in domains where comparable authorities routinely exhibit more dispersed award patterns.

These diagnostics must be interpreted within their institutional context. Framework agreements and Dynamic Purchasing Systems legitimately structure repeated transactions. Recurrence within such mechanisms may reflect institutional design rather than failure. The framework is most informative when it identifies recurrence that persists beyond what legitimate mechanisms justify.

The practical integration pathway runs through Greece's existing digital infrastructure. KIMDIS collects the data, ESIDIS administers the procedures, EADHSY provides the oversight mandate. The present framework demonstrates how an analytical layer can transform compliance data into governance intelligence using interpretable signals and authority-level market-structure indicators rather than opaque algorithmic scores. The translation from mechanism profiles to specific interventions (e.g., contestability review, dependency assessment, specification reform) should be treated as a governance hypothesis. Its practical value depends on whether these signals improve audit targeting and lead to actionable findings. A natural next step is to pilot the framework with oversight practitioners (e.g., EADHSY) by applying it to a small set of domains/authorities and evaluating whether flagged cases correspond to identifiable procedural constraints, market-structure issues, or specification practices.

### **5.3 Alternative Explanations and Limitations**

High predictability in medical domains may reflect legitimate structural constraints (certification requirements, thin markets, or centralised purchasing) rather than institutional rigidity. Similarly, low recall in construction and services may indicate genuinely competitive markets where project-specific factors override relational history. The diagnostic framework makes these distinctions visible but cannot resolve them definitively. Contextual investigation remains essential.

The temporal decline in recall across most domains admits two competing interpretations. It may reflect genuine institutional opening, a loosening of entrenched relationships as procurement reforms, digital transparency, and competitive pressures take hold. Alternatively, it may indicate growing complexity in the determinants of recurrence, with factors not captured by the three network signals (e.g., framework agreement dynamics, emergency procurement during COVID, or shifts in centralised purchasing policy) increasingly shaping buyer–supplier relationships. Distinguishing between these interpretations requires integration of procedural metadata that KIMDIS alone does not provide.

Methodological limitations include: (i) KIMDIS data quality issues (delayed entries, entity resolution challenges, occasional duplicates), (ii) absence of subcontracting, amendment, and performance data, (iii) a binary recurrence operationalisation that does not differentiate procedural types or multi-year contracts, (iv) aggregation effects from the five-digit CPV classification, (v) the three features, while interpretable, do not

exhaust available structural information and (vi) the absence of qualitative validation through practitioner engagement. The mapping from diagnostic signals to governance actions is not validated in this study. Evaluating it requires linking flagged cases to audit outcomes or qualitative case assessments.

One predictor, Historical Frequency (HF), summarises prior-year activity of the same authority–supplier pair and is therefore outcome-adjacent: high HF values mechanically imply greater persistence risk in many procurement settings. While the temporal design prevents information leakage (predictors are computed from history up to  $t$  and evaluated against outcomes in  $t+1$ ), the prominence of HF should be interpreted as a diagnostic baseline for continuity rather than as a causal mechanism, and temporal shifts in feature importance should be read with this in mind.

## 6. Conclusion

This study demonstrates that recurring buyer–supplier relationships in Greek public procurement encode measurable institutional dynamics whose governance significance varies systematically across procurement domains, over time, and among individual contracting authorities. Modelling KIMDIS award data as a temporally evolving bipartite network with three interpretable recurrence signals, the analysis identifies distinct structural configurations that correspond to different institutional logics and call for differentiated oversight responses.

Across twelve CPV domains and seven year-pairs (2018–2025), the model consistently discriminates recurring from non-recurring pairs (AUC 0.80–0.96), but the sources of that predictability differ markedly: historical continuity dominates in medical and pharmaceutical domains, structural concentration shapes repair and maintenance markets, and context-bound repetition prevails in construction and services. These distinctions imply different governance responses (contestability reviews, dependency assessments, and specification reform, respectively) rather than uniform thresholds that would produce both excessive false positives and false negatives.

The temporal dimension reveals a concurrent shift: HF dominates early year-pairs, but PA and AA contributions increase progressively, indicating that the institutional basis of recurrence diversifies over time. Authority-level analysis adds a further layer: within a single CPV category in a single year, vendor diversity ranges from 1 to over 300 and HHI from 0.008 to 1.0, underscoring that proportionate oversight must operate at the authority level rather than through aggregate domain characterisations.

The practical integration pathway runs through Greece's existing digital infrastructure (KIMDIS, ESIDIS, and EADHSY) where the missing element is the analytical layer that transforms compliance data into interpretable governance intelligence.

Several directions for future research follow directly. First, integrating procedural metadata from ESIDIS would enable decomposition of recurrence by procedure type, distinguishing framework-mediated repetition from discretionary supplier continuity.

Second, incorporating contract performance data would connect structural diagnostics to substantive public value outcomes. Third, mixed-methods research combining network diagnostics with practitioner interviews would provide the qualitative grounding needed to validate the governance significance of the structural profiles.

The broader implication is that predictability in public procurement is neither inherently beneficial nor inherently problematic. What matters for governance is whether oversight institutions possess the diagnostic capacity to distinguish among the institutional mechanisms that produce it and the interpretive frameworks to respond proportionately.

**Declaration of generative AI and AI-assisted technologies in the manuscript preparation process.**

During the preparation of this work, the authors used ChatGPT (OpenAI, GPT-5.2) to improve language clarity, grammar, and stylistic consistency, and Claude (Anthropic, Claude Sonnet 4.6) to assist in developing the Python scripts for the XGBoost classification pipeline and feature engineering. After using these tools, the authors reviewed and edited all outputs as needed and take full responsibility for the content of the published article.

## References

1. Adamic, L. A., and Adar, E. (2003). Friends and neighbors on the web. *Social Networks*, 25(3), 211–230. 10.1016/S0378-8733(03)00009-1
2. Albalade, Daniel and Bel, Germà and González-Gómez, Francisco and Picazo-Tadeo, Andrés J. (2020). Contract Renewal in Urban Water Services, Incumbent Advantage, and Market Concentration. *Public Administration Review*. 80(6), 1067-1080. 10.1111/puar.13282
3. Baltrunaite, A., Giorgiantonio, C., Mocetti, S., and Orlando, T. (2021). Discretion and supplier selection in public procurement. *Journal of Law, Economics, and Organization*, 37(1), 134–166. 10.1093/jleo/ewaa009
4. Bandiera, O., Prat, A., and Valletti, T. (2009). Active and passive waste in government spending: Evidence from a policy experiment. *American Economic Review*, 99(4), 1278–1308. 10.1257/aer.99.4.1278
5. Bannister, F., and Connolly, R. (2020). The future ain't what it used to be: Forecasting the impact of ICT on the public sphere. *Government Information Quarterly*, 37(1), 101410. 10.1016/j.giq.2019.101410
6. Barabási, A.-L., and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512. 10.1126/science.286.5439.509
7. Bozeman, B. (2002). Public-value failure: When efficient markets may not do. *Public Administration Review*, 62(2), 145–161. 10.1111/0033-3352.00165
8. Bryson, John M. and Crosby, Barbara C. and Bloomberg, Linda. (2014). Public Value Governance: Moving Beyond Traditional Public Administration and the New Public Management. *Public Administration Review*. 74(4), 445-456. 10.1111/puar.12238
9. Camboni, Riccardo and Valbonesi, Paola. (2020). Incumbent and Entrant Bidding in Scoring Rule Auctions: A Study on Italian Canteen Services. *The B.E. Journal of Economic Analysis and Policy*. 20(3). 10.1515/bejeap-2019-0146
10. Chen, T., and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). 10.1145/2939672.2939785
11. Coviello, D., and Mariniello, M. (2014). Publicity requirements in public procurement: Evidence from a regression discontinuity design. *Journal of Public Economics*, 109, 76–100. <https://doi.org/10.1016/j.jpubeco.2013.10.008>
12. Fazekas, M., and Toth, I. J. (2016). From corruption to state capture: A new analytical framework with empirical applications from Hungary. *Political Research Quarterly*, 69(2), 320-334. 10.1177/1065912916639137
13. Fontana, N., and d'Agostino, G. (2025). Threshold of shadows: Unveiling organised crime in Italian municipal public procurement. *European Journal of Political Economy*, 90, Part B. 102752. 10.1016/j.ejpoleco.2025.102752

14. Fountoukidis, I.G., Antoniou, I.E. and Varsakelis, N.C. (2023). Competitive conditions in the public procurement markets: an investigation with network analysis. *J. Ind. Bus. Econ.* 50, 347–368. 10.1007/s40812-022-00251-z
15. Fountoukidis, Ioannis and Antoniou, Ioannis and Varsakelis, Nikos. (2025). Network analysis for detecting domestic preferences in EU public procurement: enhancing market transparency and competitiveness. *International Journal of Public Sector Management.* 38. 10.1108/IJPSM-03-2024-0075.
16. Gallego, Jorge and Rivero, Gonzalo and Martinez, Juan. (2021). Preventing rather than punishing: An early warning model of malfeasance in public procurement. *International Journal of Forecasting.* 37(3), 1016-1031. 10.1016/j.ijforecast.2020.06.006.
17. Giotopoulos, I., Pavlou G., Svetoslav D., and Tsakanikas, A. (2015). Public procurement reforms in Greece: The impact of improved transparency on government expenditures. *Journal of Public Procurement*, 15(4), 458-475. 10.1108/JOPP-15-04-2015-B003
18. Herrera Fabiola, Torres Romina, Nicolis Orietta and Salas Rodrigo (2020). Characterization of the Chilean Public Procurement Ecosystem Using Social Network Analysis. *IEEE Access PP (99)*, 1-1. 10.1109/access.2020.3011947
19. Kistler, Justin T. and Sharma, Luv and Jayaram, Jayanth and Eckerd, Stephanie. (2024). Does history really repeat itself? An empirical investigation of recurring misconduct violations in public procurement. *Journal of Purchasing and Supply Management.* 29(5), 100893. 10.1016/j.pursup.2023.100893.
20. Kontogeorga, G., and Angelaras, D.N. (2023). The audit of public procurements: Evidence from Greece, in M. Tutino, V. Santolamazza, & A. Kostyuk (Eds.), *New outlooks for the scholarly research in corporate governance* (pp. 19–22). *Virtus Interpress.* 10.22495/nosrcgp3.
21. Kuo, F., and Filz, T. (2024). Unravelling the public procurement networks of architectural services in Finland as pathways of transformative innovations and tacit knowledge. *Architectural Research in Finland* 8(1):32. 10.37457/arf.146882.
22. Laffont, J.-J., and Tirole, J. (1993). *A Theory of Incentives in Procurement and Regulation.* MIT Press.
23. Levitt, Barbara and March, James G. (1988). Organizational Learning. *Annual Review of Sociology.* 14, 319-340. 10.1146/annurev.so.14.080188.001535.
24. March, James G. and Olsen, Johan P. (1989). *Rediscovering Institutions: The Organizational Basis of Politics.* Free Press, New York.
25. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., and Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data and Society*, 3(2), 1–21. 10.1177/2053951716679679.
26. Moore, M. H. (1995). *Creating Public Value: Strategic Management in Government.* Harvard University Press.

27. OECD. (2016). Fighting Bid Rigging in Public Procurement: Report on Implementing the OECD Recommendation. OECD Publishing. 10.1787/57b223af-en
28. Palguta, J., and Pertold, F. (2017). Manipulation of procurement contracts: Evidence from the introduction of discretionary thresholds. *American Economic Journal: Economic Policy*, 9(2), 293–315. 10.1257/pol.20150511
29. Peters, B. G., Pierre, J., and Randma-Liiv, T. (2011). Global financial crisis, public administration and governance: Do new problems require new solutions? *Public Organization Review*, 11(1), 13–27. 10.1007/s11115-010-0148-x
30. Pierson, P. (2000). Increasing returns, path dependence, and the study of politics. *American Political Science Review*, 94(2), 251–267. <https://doi.org/10.2307/2586011>
31. Plaček, Michal and Schmidt, Martin and Ochrana, František and Vaceková, Gabriela and Soukopová, Jana. (2019). Repeated selection of supplier, factors and results: the Czech Republic. *Journal of Public Procurement*. 19(3), 224-244. 10.1108/JOPP-12-2018-0049
32. Pliatsidis, A. (2024). Analyzing concentration in the Greek public procurement market: a network theory approach. *J. Ind. Bus. Econ.* 51, 431-480. 10.1007/s40812-023-00291-z
33. Popa, Mircea. (2019). Uncovering the structure of public procurement transactions. *Business and Politics*. 21(3), 351-383. 10.1017/bap.2019.1
34. Sturm, N.F., Candia, C., Damásio, B. et al. (2025). High earnings through firm influence: the role of hierarchical structures in public procurement. *EPJ Data Sci.* 14, 27. 10.1140/epjds/s13688-025-00543-z
35. Thai, K. V. (2009). Public Procurement: Concepts and Practices. In K. V. Thai (Ed.), *International Handbook of Public Procurement* (pp. 1–24). CRC Press, Boca Raton, FL.
36. Veale, M., and Binns, R. (2017). Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. *Big Data and Society*, 4(2), 1–17. 10.1177/2053951717743530
37. Wachs, J., Yasseri, T., Lengyel, B., and Kertész, J. (2019). Social capital predicts corruption risk in towns. *R Soc Open Sci.* 1 April 2019; 6 (4): 182103. 10.1098/rsos.182103
38. Wirtz, B. W., Weyerer, J. C., and Geyer, C. (2019). Artificial Intelligence and the Public Sector—Applications and Challenges. *International Journal of Public Administration*, 42(7), 596–615. 10.1080/01900692.2018.1498103
39. Yakovlev, Andrei and Vyglovsky, Oleg and Demidova, Olga and Bashlyk, Alexander. (2015). Incentives for Repeated Contracts in Public Procurement: Empirical Study of Gasoline Procurement in Russia. SSRN Working Paper. 10.2139/ssrn.2548461
40. Ziewitz, M. (2016). Governing algorithms: Myth, mess, and methods. *Science, Technology, and Human Values*, 41(1), 3–16. [10.1177/0162243915608948](https://doi.org/10.1177/0162243915608948)

## Appendix

The XGBoost classifier was trained using a fixed hyperparameter configuration ( $n\_estimators = 600$ ,  $max\_depth = 6$ ,  $learning\_rate = 0.05$ ,  $subsample = 0.8$ ,  $colsample\_bytree = 0.8$ ,  $objective = binary:logistic$ ) applied uniformly across all CPV categories and year-pairs. This design choice prioritises comparability of results across experimental conditions over individual model optimisation. No domain-specific or year-specific tuning was performed, ensuring that cross-domain differences in predictive performance reflect variation in the underlying network structure rather than variation in model specification. Sensitivity to data configuration is assessed through systematic variation of the negative sampling ratio (0.1, 0.5, 1.0), which alters the class balance presented to the classifier. Full results are reported in Table A1.

Table A1. Feature Importance (Gain) Stability Across Negative Sampling Ratios

| Category |                            | HF Gain |       |       | PA Gain |       |       | AA Gain |       |       | Dominant Feature |        |          |
|----------|----------------------------|---------|-------|-------|---------|-------|-------|---------|-------|-------|------------------|--------|----------|
| CPV      | Category                   | 0.1     | 0.5   | 1.0   | 0.1     | 0.5   | 1.0   | 0.1     | 0.5   | 1.0   | at 0.1           | at 1.0 | Stable ? |
| 30192    | Office Supplies            | 0.535   | 0.563 | 0.570 | 0.210   | 0.158 | 0.131 | 0.254   | 0.279 | 0.299 | HF               | HF     | ✓        |
| 44111    | Building Materials         | 0.508   | 0.553 | 0.556 | 0.201   | 0.156 | 0.134 | 0.291   | 0.291 | 0.310 | HF               | HF     | ✓        |
| 45233    | Road Construction/Repair   | 0.458   | 0.498 | 0.534 | 0.245   | 0.179 | 0.141 | 0.296   | 0.322 | 0.325 | HF               | HF     | ✓        |
| 50112    | Motor Vehicle Maintenance  | 0.446   | 0.532 | 0.534 | 0.285   | 0.200 | 0.175 | 0.269   | 0.268 | 0.291 | HF               | HF     | ✓        |
| 90910    | Cleaning Services          | 0.471   | 0.500 | 0.533 | 0.262   | 0.196 | 0.160 | 0.267   | 0.304 | 0.308 | HF               | HF     | ✓        |
| 50000    | Repair & Maintenance       | 0.490   | 0.530 | 0.520 | 0.252   | 0.190 | 0.162 | 0.258   | 0.280 | 0.318 | HF               | HF     | ✓        |
| 90911    | Building Cleaning Services | 0.453   | 0.511 | 0.511 | 0.269   | 0.196 | 0.172 | 0.278   | 0.293 | 0.318 | HF               | HF     | ✓        |
| 79713    | Guard & Security Services  | 0.397   | 0.457 | 0.467 | 0.300   | 0.225 | 0.186 | 0.303   | 0.318 | 0.347 | HF               | HF     | ✓        |
| 33696    | Reagents & Lab Chemicals   | 0.390   | 0.421 | 0.447 | 0.244   | 0.211 | 0.188 | 0.366   | 0.369 | 0.365 | HF               | HF     | ✓        |
| 33141    | Minor Medical Devices      | 0.353   | 0.397 | 0.417 | 0.279   | 0.252 | 0.234 | 0.369   | 0.351 | 0.350 | AA               | HF     | ✗        |
| 33140    | Medical Consumables        | 0.309   | 0.364 | 0.386 | 0.288   | 0.247 | 0.223 | 0.402   | 0.389 | 0.391 | AA               | AA     | ✓        |
| 33111    | Radiography Equipment      | 0.292   | 0.382 | 0.356 | 0.305   | 0.258 | 0.263 | 0.403   | 0.360 | 0.381 | AA               | AA     | ✓        |

Hyperparameters were fixed across domains to ensure comparability of performance and mechanism decomposition. Because gain-based feature importance can vary with model complexity, we conducted a targeted sensitivity analysis on a representative subset of domains and year-pairs, varying  $max\_depth$  and  $learning\_rate$  over a small grid. AUC values remain stable and the qualitative mechanism profiles are largely preserved; occasional swaps between HF and AA occur only in cases where their gain shares are near-tied.

Table A2. Hyperparameter sensitivity of gain-based mechanism decomposition

| CPV   | Period            | Baseline dominant | Baseline HF/PA/AA     | AUC (baseline) | HF (min-max) | PA (min-max) | AA (min-max) | AUC (min-max) | Dominant stable? | Mean $\rho$ vs baseline |
|-------|-------------------|-------------------|-----------------------|----------------|--------------|--------------|--------------|---------------|------------------|-------------------------|
| 33111 | early (2018→2019) | HF                | 0.516 / 0.167 / 0.317 | 0.838          | 0.428–0.567  | 0.133–0.240  | 0.297–0.348  | 0.810–0.881   | Yes              | 1.000                   |
| 33111 | late (2024→2025)  | HF                | 0.429 / 0.199 / 0.372 | 0.954          | 0.404–0.504  | 0.163–0.208  | 0.291–0.418  | 0.950–0.959   | Mostly (HF)      | 0.944                   |
| 33696 | early (2018→2019) | HF                | 0.796 / 0.070 / 0.134 | 0.949          | 0.606–0.805  | 0.059–0.150  | 0.134–0.245  | 0.944–0.953   | Yes              | 1.000                   |
| 33696 | late (2024→2025)  | HF                | 0.469 / 0.158 / 0.373 | 0.943          | 0.438–0.520  | 0.120–0.214  | 0.321–0.383  | 0.932–0.950   | Yes              | 1.000                   |
| 50000 | early (2018→2019) | HF                | 0.951 / 0.016 / 0.033 | 0.880          | 0.891–0.974  | 0.009–0.032  | 0.017–0.077  | 0.876–0.885   | Yes              | 1.000                   |
| 50000 | late (2024→2025)  | HF                | 0.415 / 0.190 / 0.395 | 0.877          | 0.392–0.470  | 0.125–0.243  | 0.340–0.426  | 0.858–0.879   | Mostly (HF)      | 0.889                   |

Notes: Baseline hyperparameters: max\_depth = 6, learning\_rate = 0.05. Alternative settings vary max\_depth  $\in$  {4, 6, 8} and learning\_rate  $\in$  {0.03, 0.05, 0.10}, holding the data, sampling, and evaluation procedure fixed. HF/PA/AA entries are normalised gain shares.  $\rho$  is Spearman correlation of the 3-element gain-share vector [HF, PA, AA] relative to baseline.

Table A3: Full Predictive Performance by CPV Category and Year-Pair (Negative Ratio = 1.0)

| CPV   | Category              | Year-Pair | N Pos | N Neg | Eff. Ratio | AUC   | Precision | Recall | F1    | HF Gain | PA Gain | AA Gain | Ceiling |
|-------|-----------------------|-----------|-------|-------|------------|-------|-----------|--------|-------|---------|---------|---------|---------|
| 30192 | Office Supplies       | 2018→2019 | 2658  | 2658  | 1          | 0.866 | 0.561     | 0.223  | 0.319 | 0.924   | 0.019   | 0.057   | No      |
| 30192 | Office Supplies       | 2019→2020 | 4305  | 4305  | 1          | 0.864 | 0.510     | 0.252  | 0.337 | 0.613   | 0.105   | 0.283   | No      |
| 30192 | Office Supplies       | 2020→2021 | 5551  | 5551  | 1          | 0.866 | 0.562     | 0.244  | 0.340 | 0.494   | 0.146   | 0.360   | No      |
| 30192 | Office Supplies       | 2021→2022 | 6897  | 6897  | 1          | 0.885 | 0.528     | 0.199  | 0.289 | 0.577   | 0.141   | 0.283   | No      |
| 30192 | Office Supplies       | 2022→2023 | 7886  | 7886  | 1          | 0.875 | 0.510     | 0.204  | 0.291 | 0.495   | 0.140   | 0.365   | No      |
| 30192 | Office Supplies       | 2023→2024 | 8965  | 8965  | 1          | 0.860 | 0.450     | 0.157  | 0.233 | 0.447   | 0.190   | 0.363   | No      |
| 30192 | Office Supplies       | 2024→2025 | 9801  | 9801  | 1          | 0.866 | 0.446     | 0.127  | 0.198 | 0.439   | 0.180   | 0.382   | No      |
| 33111 | Radiography Equipment | 2018→2019 | 268   | 268   | 1          | 0.831 | 0.625     | 0.441  | 0.517 | 0.431   | 0.199   | 0.371   | No      |
| 33111 | Radiography Equipment | 2019→2020 | 447   | 447   | 1          | 0.949 | 0.790     | 0.714  | 0.750 | 0.204   | 0.350   | 0.446   | No      |
| 33111 | Radiography Equipment | 2020→2021 | 589   | 589   | 1          | 0.921 | 0.611     | 0.478  | 0.537 | 0.235   | 0.359   | 0.407   | No      |
| 33111 | Radiography Equipment | 2021→2022 | 739   | 739   | 1          | 0.912 | 0.594     | 0.396  | 0.475 | 0.362   | 0.285   | 0.353   | No      |
| 33111 | Radiography Equipment | 2022→2023 | 1351  | 1351  | 1          | 0.906 | 0.439     | 0.333  | 0.379 | 0.429   | 0.218   | 0.353   | No      |
| 33111 | Radiography Equipment | 2023→2024 | 2052  | 2052  | 1          | 0.944 | 0.580     | 0.403  | 0.475 | 0.402   | 0.234   | 0.365   | No      |
| 33111 | Radiography Equipment | 2024→2025 | 2671  | 2671  | 1          | 0.952 | 0.463     | 0.329  | 0.385 | 0.429   | 0.201   | 0.371   | No      |
| 33140 | Medical Consumables   | 2018→2019 | 2611  | 2611  | 1          | 0.925 | 0.660     | 0.543  | 0.595 | 0.717   | 0.091   | 0.192   | No      |
| 33140 | Medical Consumables   | 2019→2020 | 4431  | 4431  | 1          | 0.944 | 0.677     | 0.613  | 0.644 | 0.240   | 0.362   | 0.398   | No      |

|       |                          |           |       |       |   |       |       |       |       |       |       |       |    |
|-------|--------------------------|-----------|-------|-------|---|-------|-------|-------|-------|-------|-------|-------|----|
| 33140 | Medical Consumables      | 2020→2021 | 6822  | 6822  | 1 | 0.940 | 0.660 | 0.497 | 0.567 | 0.259 | 0.298 | 0.444 | No |
| 33140 | Medical Consumables      | 2021→2022 | 8395  | 8395  | 1 | 0.932 | 0.646 | 0.362 | 0.464 | 0.417 | 0.199 | 0.384 | No |
| 33140 | Medical Consumables      | 2022→2023 | 9766  | 9766  | 1 | 0.934 | 0.621 | 0.379 | 0.471 | 0.374 | 0.207 | 0.419 | No |
| 33140 | Medical Consumables      | 2023→2024 | 11332 | 11332 | 1 | 0.936 | 0.616 | 0.379 | 0.469 | 0.351 | 0.211 | 0.438 | No |
| 33140 | Medical Consumables      | 2024→2025 | 12521 | 12521 | 1 | 0.934 | 0.668 | 0.398 | 0.499 | 0.341 | 0.197 | 0.461 | No |
| 33141 | Minor Medical Devices    | 2018→2019 | 2492  | 2492  | 1 | 0.929 | 0.654 | 0.634 | 0.644 | 0.694 | 0.120 | 0.186 | No |
| 33141 | Minor Medical Devices    | 2019→2020 | 4288  | 4288  | 1 | 0.928 | 0.679 | 0.466 | 0.553 | 0.277 | 0.321 | 0.403 | No |
| 33141 | Minor Medical Devices    | 2020→2021 | 5974  | 5974  | 1 | 0.923 | 0.620 | 0.362 | 0.457 | 0.402 | 0.278 | 0.320 | No |
| 33141 | Minor Medical Devices    | 2021→2022 | 7429  | 7429  | 1 | 0.928 | 0.616 | 0.340 | 0.438 | 0.402 | 0.227 | 0.372 | No |
| 33141 | Minor Medical Devices    | 2022→2023 | 8830  | 8830  | 1 | 0.914 | 0.524 | 0.272 | 0.358 | 0.388 | 0.240 | 0.371 | No |
| 33141 | Minor Medical Devices    | 2023→2024 | 10196 | 10196 | 1 | 0.918 | 0.581 | 0.306 | 0.401 | 0.387 | 0.223 | 0.390 | No |
| 33141 | Minor Medical Devices    | 2024→2025 | 11273 | 11273 | 1 | 0.927 | 0.575 | 0.308 | 0.401 | 0.366 | 0.227 | 0.407 | No |
| 33696 | Reagents & Lab Chemicals | 2018→2019 | 1493  | 1493  | 1 | 0.947 | 0.786 | 0.710 | 0.746 | 0.766 | 0.085 | 0.149 | No |
| 33696 | Reagents & Lab Chemicals | 2019→2020 | 2369  | 2369  | 1 | 0.932 | 0.696 | 0.689 | 0.692 | 0.278 | 0.283 | 0.439 | No |
| 33696 | Reagents & Lab Chemicals | 2020→2021 | 3226  | 3226  | 1 | 0.941 | 0.746 | 0.641 | 0.690 | 0.412 | 0.188 | 0.399 | No |
| 33696 | Reagents & Lab Chemicals | 2021→2022 | 4034  | 4034  | 1 | 0.936 | 0.682 | 0.580 | 0.627 | 0.368 | 0.234 | 0.398 | No |
| 33696 | Reagents & Lab Chemicals | 2022→2023 | 4775  | 4775  | 1 | 0.936 | 0.642 | 0.530 | 0.581 | 0.434 | 0.171 | 0.395 | No |
| 33696 | Reagents & Lab Chemicals | 2023→2024 | 5430  | 5430  | 1 | 0.942 | 0.656 | 0.475 | 0.551 | 0.407 | 0.198 | 0.396 | No |
| 33696 | Reagents & Lab Chemicals | 2024→2025 | 5848  | 5848  | 1 | 0.943 | 0.726 | 0.546 | 0.624 | 0.461 | 0.161 | 0.378 | No |
| 44111 | Building Materials       | 2018→2019 | 1358  | 1358  | 1 | 0.864 | 0.540 | 0.353 | 0.427 | 0.954 | 0.011 | 0.035 | No |
| 44111 | Building Materials       | 2019→2020 | 2243  | 2243  | 1 | 0.837 | 0.442 | 0.259 | 0.326 | 0.541 | 0.129 | 0.330 | No |
| 44111 | Building Materials       | 2020→2021 | 3062  | 3062  | 1 | 0.866 | 0.470 | 0.213 | 0.293 | 0.493 | 0.144 | 0.363 | No |
| 44111 | Building Materials       | 2021→2022 | 3772  | 3772  | 1 | 0.862 | 0.521 | 0.257 | 0.344 | 0.543 | 0.139 | 0.318 | No |
| 44111 | Building Materials       | 2022→2023 | 4392  | 4392  | 1 | 0.886 | 0.655 | 0.356 | 0.461 | 0.509 | 0.142 | 0.349 | No |
| 44111 | Building Materials       | 2023→2024 | 5050  | 5050  | 1 | 0.880 | 0.577 | 0.262 | 0.360 | 0.434 | 0.189 | 0.377 | No |
| 44111 | Building Materials       | 2024→2025 | 5571  | 5571  | 1 | 0.867 | 0.440 | 0.179 | 0.255 | 0.419 | 0.182 | 0.399 | No |
| 45233 | Road Construction/Repair | 2018→2019 | 2145  | 2145  | 1 | 0.868 | 0.511 | 0.411 | 0.455 | 0.947 | 0.010 | 0.043 | No |
| 45233 | Road Construction/Repair | 2019→2020 | 3576  | 3576  | 1 | 0.856 | 0.500 | 0.253 | 0.336 | 0.532 | 0.129 | 0.339 | No |
| 45233 | Road Construction/Repair | 2020→2021 | 4561  | 4561  | 1 | 0.884 | 0.586 | 0.297 | 0.394 | 0.487 | 0.143 | 0.370 | No |
| 45233 | Road Construction/Repair | 2021→2022 | 5581  | 5581  | 1 | 0.858 | 0.526 | 0.223 | 0.313 | 0.434 | 0.154 | 0.413 | No |
| 45233 | Road Construction/Repair | 2022→2023 | 6708  | 6708  | 1 | 0.878 | 0.488 | 0.174 | 0.256 | 0.516 | 0.148 | 0.337 | No |
| 45233 | Road Construction/Repair | 2023→2024 | 7895  | 7895  | 1 | 0.856 | 0.302 | 0.081 | 0.128 | 0.386 | 0.210 | 0.405 | No |
| 45233 | Road Construction/Repair | 2024→2025 | 8606  | 8606  | 1 | 0.854 | 0.500 | 0.113 | 0.185 | 0.438 | 0.190 | 0.372 | No |

|       |                            |           |      |      |   |       |       |       |       |       |       |       |    |
|-------|----------------------------|-----------|------|------|---|-------|-------|-------|-------|-------|-------|-------|----|
| 50000 | Repair & Maintenance       | 2018→2019 | 2130 | 2130 | 1 | 0.880 | 0.642 | 0.274 | 0.384 | 0.951 | 0.016 | 0.033 | No |
| 50000 | Repair & Maintenance       | 2019→2020 | 3752 | 3752 | 1 | 0.889 | 0.598 | 0.275 | 0.377 | 0.529 | 0.163 | 0.309 | No |
| 50000 | Repair & Maintenance       | 2020→2021 | 5261 | 5261 | 1 | 0.864 | 0.492 | 0.148 | 0.228 | 0.441 | 0.189 | 0.371 | No |
| 50000 | Repair & Maintenance       | 2021→2022 | 6502 | 6502 | 1 | 0.867 | 0.429 | 0.138 | 0.209 | 0.450 | 0.182 | 0.368 | No |
| 50000 | Repair & Maintenance       | 2022→2023 | 7692 | 7692 | 1 | 0.869 | 0.467 | 0.133 | 0.207 | 0.419 | 0.200 | 0.381 | No |
| 50000 | Repair & Maintenance       | 2023→2024 | 8764 | 8764 | 1 | 0.874 | 0.278 | 0.096 | 0.143 | 0.431 | 0.196 | 0.373 | No |
| 50000 | Repair & Maintenance       | 2024→2025 | 9642 | 9642 | 1 | 0.877 | 0.415 | 0.111 | 0.175 | 0.417 | 0.190 | 0.393 | No |
| 50112 | Motor Vehicle Maintenance  | 2018→2019 | 1475 | 1475 | 1 | 0.941 | 0.745 | 0.577 | 0.650 | 0.945 | 0.021 | 0.034 | No |
| 50112 | Motor Vehicle Maintenance  | 2019→2020 | 2524 | 2524 | 1 | 0.929 | 0.642 | 0.417 | 0.505 | 0.479 | 0.227 | 0.294 | No |
| 50112 | Motor Vehicle Maintenance  | 2020→2021 | 3399 | 3399 | 1 | 0.913 | 0.616 | 0.393 | 0.480 | 0.445 | 0.213 | 0.342 | No |
| 50112 | Motor Vehicle Maintenance  | 2021→2022 | 4185 | 4185 | 1 | 0.897 | 0.575 | 0.302 | 0.396 | 0.466 | 0.195 | 0.339 | No |
| 50112 | Motor Vehicle Maintenance  | 2022→2023 | 4911 | 4911 | 1 | 0.897 | 0.568 | 0.295 | 0.389 | 0.469 | 0.188 | 0.344 | No |
| 50112 | Motor Vehicle Maintenance  | 2023→2024 | 5622 | 5622 | 1 | 0.894 | 0.546 | 0.263 | 0.355 | 0.492 | 0.181 | 0.327 | No |
| 50112 | Motor Vehicle Maintenance  | 2024→2025 | 6249 | 6249 | 1 | 0.881 | 0.509 | 0.209 | 0.297 | 0.444 | 0.202 | 0.355 | No |
| 79713 | Guard & Security Services  | 2018→2019 | 484  | 484  | 1 | 0.823 | 0.577 | 0.250 | 0.349 | 0.933 | 0.016 | 0.051 | No |
| 79713 | Guard & Security Services  | 2019→2020 | 737  | 737  | 1 | 0.864 | 0.487 | 0.279 | 0.355 | 0.422 | 0.189 | 0.388 | No |
| 79713 | Guard & Security Services  | 2020→2021 | 1071 | 1071 | 1 | 0.872 | 0.508 | 0.313 | 0.388 | 0.375 | 0.220 | 0.405 | No |
| 79713 | Guard & Security Services  | 2021→2022 | 1434 | 1434 | 1 | 0.852 | 0.507 | 0.282 | 0.362 | 0.419 | 0.202 | 0.379 | No |
| 79713 | Guard & Security Services  | 2022→2023 | 1737 | 1737 | 1 | 0.832 | 0.500 | 0.282 | 0.361 | 0.383 | 0.211 | 0.406 | No |
| 79713 | Guard & Security Services  | 2023→2024 | 2015 | 2015 | 1 | 0.842 | 0.364 | 0.214 | 0.270 | 0.384 | 0.234 | 0.383 | No |
| 79713 | Guard & Security Services  | 2024→2025 | 2289 | 2289 | 1 | 0.830 | 0.440 | 0.173 | 0.249 | 0.354 | 0.229 | 0.417 | No |
| 90910 | Cleaning Services          | 2018→2019 | 1095 | 1095 | 1 | 0.851 | 0.563 | 0.164 | 0.254 | 0.945 | 0.015 | 0.041 | No |
| 90910 | Cleaning Services          | 2019→2020 | 1745 | 1745 | 1 | 0.828 | 0.397 | 0.218 | 0.281 | 0.575 | 0.140 | 0.286 | No |
| 90910 | Cleaning Services          | 2020→2021 | 2445 | 2445 | 1 | 0.880 | 0.563 | 0.250 | 0.346 | 0.466 | 0.181 | 0.354 | No |
| 90910 | Cleaning Services          | 2021→2022 | 2970 | 2970 | 1 | 0.865 | 0.493 | 0.232 | 0.316 | 0.520 | 0.174 | 0.306 | No |
| 90910 | Cleaning Services          | 2022→2023 | 3410 | 3410 | 1 | 0.869 | 0.514 | 0.209 | 0.297 | 0.383 | 0.195 | 0.422 | No |
| 90910 | Cleaning Services          | 2023→2024 | 3908 | 3908 | 1 | 0.860 | 0.438 | 0.187 | 0.262 | 0.454 | 0.208 | 0.339 | No |
| 90910 | Cleaning Services          | 2024→2025 | 4394 | 4394 | 1 | 0.860 | 0.563 | 0.207 | 0.303 | 0.388 | 0.206 | 0.406 | No |
| 90911 | Building Cleaning Services | 2018→2019 | 659  | 659  | 1 | 0.876 | 0.667 | 0.354 | 0.463 | 0.910 | 0.022 | 0.068 | No |
| 90911 | Building Cleaning Services | 2019→2020 | 1075 | 1075 | 1 | 0.861 | 0.493 | 0.379 | 0.429 | 0.466 | 0.182 | 0.353 | No |
| 90911 | Building Cleaning Services | 2020→2021 | 1553 | 1553 | 1 | 0.834 | 0.475 | 0.252 | 0.329 | 0.424 | 0.195 | 0.381 | No |
| 90911 | Building Cleaning Services | 2021→2022 | 1891 | 1891 | 1 | 0.869 | 0.559 | 0.295 | 0.386 | 0.431 | 0.186 | 0.383 | No |
| 90911 | Building Cleaning Services | 2022→2023 | 2268 | 2268 | 1 | 0.838 | 0.384 | 0.204 | 0.267 | 0.424 | 0.209 | 0.367 | No |

|       |                            |           |      |      |   |       |       |       |       |       |       |       |    |
|-------|----------------------------|-----------|------|------|---|-------|-------|-------|-------|-------|-------|-------|----|
| 90911 | Building Cleaning Services | 2023→2024 | 2597 | 2597 | 1 | 0.875 | 0.513 | 0.285 | 0.366 | 0.504 | 0.196 | 0.300 | No |
| 90911 | Building Cleaning Services | 2024→2025 | 2910 | 2910 | 1 | 0.845 | 0.406 | 0.182 | 0.251 | 0.415 | 0.214 | 0.372 | No |

Table A4. Robustness of Predictive Performance Across Negative Sampling Ratios

| Category |                            | Ratio = 0.1 |       |       |      | Ratio = 0.5 |       |       |      | Ratio = 1.0 |       |       |      | Rank Shift (0.1→1.0) |
|----------|----------------------------|-------------|-------|-------|------|-------------|-------|-------|------|-------------|-------|-------|------|----------------------|
| CPV      | Category                   | AUC         | SD    | F1    | Rank | AUC         | SD    | F1    | Rank | AUC         | SD    | F1    | Rank | ΔRank                |
| 33696    | Reagents & Lab Chemicals   | 0.877       | 0.007 | 0.653 | 1    | 0.919       | 0.008 | 0.656 |      | 0.939       | 0.005 | 0.644 | 1    | 0                    |
| 33140    | Medical Consumables        | 0.874       | 0.008 | 0.539 | 2    | 0.911       | 0.009 | 0.540 |      | 0.935       | 0.006 | 0.530 | 2    | 0                    |
| 33141    | Minor Medical Devices      | 0.847       | 0.012 | 0.484 | 4    | 0.892       | 0.007 | 0.446 |      | 0.924       | 0.006 | 0.464 | 3    | 1                    |
| 33111    | Radiography Equipment      | 0.849       | 0.075 | 0.527 | 3    | 0.889       | 0.053 | 0.520 |      | 0.916       | 0.042 | 0.503 | 4    | 1                    |
| 50112    | Motor Vehicle Maintenance  | 0.803       | 0.037 | 0.426 | 5    | 0.865       | 0.020 | 0.435 |      | 0.907       | 0.021 | 0.439 | 5    | 0                    |
| 50000    | Repair & Maintenance       | 0.752       | 0.013 | 0.233 | 6    | 0.817       | 0.014 | 0.237 |      | 0.874       | 0.009 | 0.246 | 6    | 0                    |
| 30192    | Office Supplies            | 0.736       | 0.018 | 0.274 | 8    | 0.810       | 0.019 | 0.272 |      | 0.869       | 0.008 | 0.287 | 7    | 1                    |
| 44111    | Building Materials         | 0.739       | 0.035 | 0.321 | 7    | 0.819       | 0.031 | 0.350 |      | 0.866       | 0.015 | 0.352 | 8    | 1                    |
| 45233    | Road Construction/Repair   | 0.736       | 0.017 | 0.311 | 8    | 0.819       | 0.012 | 0.304 |      | 0.865       | 0.012 | 0.295 | 9    | 1                    |
| 90910    | Cleaning Services          | 0.712       | 0.033 | 0.324 | 10   | 0.803       | 0.018 | 0.287 |      | 0.859       | 0.016 | 0.294 | 10   | 0                    |
| 90911    | Building Cleaning Services | 0.698       | 0.044 | 0.346 | 11   | 0.789       | 0.021 | 0.344 |      | 0.857       | 0.018 | 0.356 | 11   | 0                    |
| 79713    | Guard & Security Services  | 0.671       | 0.043 | 0.333 | 12   | 0.791       | 0.016 | 0.358 |      | 0.845       | 0.018 | 0.333 | 12   | 0                    |

Note: Spearman rank correlation of mean AUC between ratio=0.1 and ratio=1.0:  $\rho = 0.9807$ . Green = stable rank. yellow = shift  $\leq 2$ . red = shift  $> 2$ . One observation excluded (CPV 90911, 2021) due to single-class outcome at ratio=0.5.

## Previous Papers in this Series

218. **Theodoros Zachariadis and Elias Giannakis**, *Slow but Tangible Progress: The Economic Dividends of Decarbonisation in Cyprus*, April 2026
217. **Christos Papagiannis, Stefanos Tyros and Sofia Vasilopoulou**, *Beyond Left and Right: Quantitative Evidence of the Establishment–Anti-Establishment Dimension in Greece*, March 2026
216. **Kyriakos Andreou, Andreas Fousteris, Sotirios Kokas, Alexandros Kontonikas and Emmanouil Pyrgiotakis**, *Modelling Greek Firms’ Survival Rates and Identifying “Zombies”*, February 2026
215. **Stelios Giannoulakis, Angelos Kanas, Marina-Eliza Spaliara and John Tsoukalas**, *U.S. Tariffs and Greek Exports*, January 2026
214. **Antonis Pastellopoulos**, *Cypriot Identity as an Alternative: Grassroots Perspectives on Cypriotness and Nationalism*, December 2025
213. **Charis Psaltis, Neophytos Loizides, Andreas Michael, Nikandros Ioannidis, Edward Morgan Jones and Laura Sudulich**, *Natural resources co-management, green transition and divided societies - Zones of agreement in the Cyprus case using a conjoint survey experiment*, November 2025
212. **Michalis Nikiforos, Vlassis Missos, Christos Pierros and Nikolaos Rodousakis**, *The café economy: Structural transformation in Greece in the wake of austerity and “reforms”*, October 2025
211. **Andreas Psarras, Theodore Panagiotidis and Andreas Andronikidis**, *Fuel price effects on motor vehicle collisions: Evidence from Greece*, September 2025
210. **Yani Kartalis**, *From Ostensible to Actual Media Pluralism. An examination of content diversity in Greece's fragmented online media system.*, August 2025
209. **Aikaterini E. Karadimitropoulou, Tryfonas Christou, Michael Chletsos and Alexandros P. Bechlioulis**, *Wage Reforms and Equality Gains: Evidence from Greece*, July 2025
208. **Theodoros Zachariadis, Constantinos Taliotis, Melina Moleskis and Pantelis Solomou**, *A Climate Neutrality Strategy for Cyprus*, June 2025
207. **Kyriakos Drivas and Afroditi Anagnosti**, *On the Role of Innovation in the Generation of Value-Added Trade Opportunities*, April 2025
206. **Athanasia Chalari**, *Transitioning from European Citizenship towards Immigration Identities after Brexit (the case of Greek diaspora in the UK)*, March 2025
205. **Manussos Marangudakis, Theodore Chadjipadelis, Maria Zampoka and Christos Kouroutzas**, *The Anti-vaccination Movement in Greece*, February 2025
204. **Alexandros Kontonikas and Emmanouil Pyrgiotakis**, *A comprehensive analysis of transactions in the Greek residential property market*, January 2025
203. **Athanasia Chalari**, *Cultural identities among Greek Diaspora in the UK*, December 2024