

Department of Economics

Athens University of Economics and Business

WORKING PAPER no. 14-2025

Measuring Textual Cohesion via Graph Curvature, Weighted Dominance, and Embedding-Based Coherence

Stelios Arvanitis

August 2025

The Working Papers in this series circulate mainly for early presentation and discussion, as well as for the information of the Academic Community and all interested in our current research activity.

The authors assume full responsibility for the accuracy of their paper as well as for the opinions expressed therein.

Measuring Textual Cohesion via Graph Curvature, Weighted Dominance, and Embedding-Based Coherence

Stelios Arvanitis
Department of Economics
Athens University of Economics and Business, Greece
stelios@aueb.gr

August 27, 2025

Abstract

Evaluating textual cohesion is an important component of information retrieval, summarization, and automated content quality assessment. We present a framework that combines semantic embeddings with graph-based indicators of discourse structure. Sentence similarity graphs are built from SBERT embeddings, and Ollivier–Ricci curvature is applied to capture structural connectivity. We introduce a Weighted Utopia Index (UI), derived from a dominance relation among curvature embeddings, to quantify proximity to cohesive discourse. To enhance robustness, a Nash-style multiobjective aggregation balances the UI with embedding-based coherence scores. Experiments on the publicly available Cohesentia dataset show that the proposed metrics achieve higher alignment with human coherence judgments than embedding-only baselines, with consistent gains in agreement and nDCG. The approach is domain-independent, interpretable, and adaptable to a range of evaluation scenarios, including summarization quality control and discourse analysis.

Keywords: Textual cohesion, Discourse structure, Semantic similarity, Ollivier–Ricci curvature, Computational discourse analysis, Weighted Utopia Index, Nash-type Score, automated text evaluation.

1 Introduction

Coherence is a fundamental property of well-structured texts, referring to the perceived logical and semantic consistency among their constituent units. In the pre-neural-networks era, coherence modeling drew heavily from linguistic theories. Local models such as the entity grid (e.g. Barzilay and Lapata (2008)), inspired by Centering Theory (Grosz et al. (1995)), focused on tracking syntactic roles of discourse entities across adjacent sentences. Discourse-based approaches leveraged frameworks like Rhetorical Structure Theory (see Mann and Thompson (1988)) to model coherence via annotated discourse relations (see for example Lin

et al. (2011)). Additional paradigms used search alogorithms on ISD-graphs, for example Soricut and Marcu (2006), topic distributions, for example Barzilay and Lee (2004), or syntactic alignment, for example Louis and Nenkova (2012), while evaluation was typically based on proxy tasks such as sentence reordering or essay scoring via regression analysis and cross validation (see Burstein et al. (1998)).

Recent work in the neural-networks era has shifted toward embedding-based methods, particularly those using transformer models to learn representations of sentence plausibility. Efforts such as Maimon and Tsarfaty (2023) address coherence modeling directly, seeking to disentangle it from related but distinct tasks such as fluency or entailment.

In parallel, graph-theoretic and information-theoretic approaches have emerged as promising complements for capturing structural patterns in text, aligning naturally with methods from statistical physics and network science. These models represent sentences or concepts as nodes and define semantic or syntactic relations as weighted edges, enabling analysis of textual coherence via network topology. An indicative and non-exhausive list of applications includes the quantification of narrative complexity for interactive systems, see Kwon et al. (2016), the use of complex networks analytics to evaluate consistency in the use of words, see Amancio et al. (2012), or the use of Shannon and Kolmogorov entropy to distinguish between ordered and stochastic components in natural written texts, see Papadimitriou et al. (2010) or perform text mining, see Jamaati and Mehri (2018). Such methods offer a complementary perspective to deep learning, grounding text structure in geometric or topological formalisms.

Building on this line of work, we introduce a new methodology that integrates semantic embeddings, the graph-theoretic Olivier-Ricci curvature, and ordering relations between the resulting curvature profiles, into a framework for quantifying rhetorical cohesion. Our primary construction lies in defining a Weighted Utopia Index (wUI), which compares the curvature profile of a given text against corpus-level upper and lower envelopes, under a trainable weight function over "rhetorical time". The weight function is optimized using human coherence ratings via ordered probit regression, allowing the model to learn cognitively aligned temporal emphasis. We benchmark this metric on the human coherence annotations endowed Cohesentia, and compare it against an SBERT embeddings-based score constructed as an average of cosine similarities between adjacent sentences, and a hybrid alternative; the latter is constructed as a multiobjective-based Nash-type score. We apply it in a case study of rhetorical strategy across politically high-profile inaugural addresses.

Structure of the Paper. Section 2 introduces the full methodological pipeline, including the construction of SBERT-based semantic graphs, computation of Ricci curvature, and definition of the Weighted UI with its training via ordinal probit estimation on the Cohesentia dataset. Section 3 presents the benchmarking experiment, detailing the Cohesentia dataset, the learning procedure for the weight functions, and quantitative evaluations of the Weighted UI compared to alternative metrics. Section 4 applies the trained methodology to an indicative corpus of inaugural US presidential speeches. Finally, Section 5 discusses extensions to the methodology and potential future applications.

2 Methodology

For a finite corpus of texts $\mathcal{D} = \{d_1, \ldots, d_N\}$, where each text d_k consists of a non-empty ordered sequence of sentences $\mathcal{S}_k = \{s_1^{(k)}, \ldots, s_{n_k}^{(k)}\}$, the goal is to construct numerical scores that approximate the internal cohesion of each d_k, by analyzing semantic and geometric information over $\mathcal{S}_k$.

The current section develops the methodological pipeline. We begin by constructing semantic similarity graphs using SBERT-based sentence embeddings (see Reimers and Gurevych (2019)), linking sentences via cosine similarity to form weighted, undirected graphs in the spirit of TextRank (Mihalcea and Tarau (2004)). The graph-theoretic Ollivier-Ricci curvature (see Ollivier (2007, 2009)) is then computed on these graphs using optimal transport distances between neighborhood distributions, yielding sentence-level curvature profiles, represented by text-by-text comparable step function embeddings of the curvature vectors.

To assess structural alignment using information from the curvature profiles, we introduce a "reverse First order Dominance-type" order relation and produce a Weighted Utopia Index (wUI). The latter is an adaptation of the Utopia Index associated with Second Order Stochastic Dominance and introduced in Anderson et al. (2020). Given the set of curvature profiles for the corpus, the wUI quantifies the severity of the obstacles to order maximality of a speech's curvature profile against corpus-wide bounds. The wUI is computed using a learned weight function $w_\theta(x) = \exp(\theta_0 + \theta_1 x + \theta_2 x^2)$, where $\theta \in \mathbb{R}^3$ is trained via ordinal probit likelihood using human coherence judgments from the Cohesentia dataset (Maimon and Tsarfaty (2023)). This endows the wUI with sensitivity to "temporal" structure-favoring coherence patterns aligned with human perception, as represented in this dataset.

To complement structural metrics, we compute the SBERT-based Coherence Score (SBERT-CS), which measures semantic continuity as the average cosine similarity between adjacent SBERT-sentence embeddings. This average partially captures local flow and topical progression across the discourse, leveraging contextual sentence representations produced by a pre-trained transformer.

Finally, the Nash Cohesion Score is defined as the geometric mean of the SBERT-based coherence score and the weighted Utopia Index. It summarizes a trade-off between semantic plausibility and structural alignment. This can be thought of as a Pareto-scalarization of the two-dimensional score obtained from vectorising SBERT-CS with wUI.

2.1 TextRank and SBERT Semantic Graph

Let $\mathcal{S}_k = \{s_1^{(k)}, \ldots, s_{n_k}^{(k)}\}$ be a finite set of sentences extracted from an arbitrary document in the corpus. The goal of the TextRank algorithm (Mihalcea and Tarau (2004)) is to compute a score $r_i^{(k)}$ for each sentence $s_i^{(k)}$, reflecting its importance in terms of semantic similarity to other sentences.

The relationships between sentences are represented using an undirected weighted graph $G^{(k)} = (V^{(k)}, E^{(k)}, W^{(k)})$, where:

- $V^{(k)} = \mathcal{S}_k$ is the set of nodes (sentences),
- $(s_i^{(k)}, s_j^{(k)}) \in E^{(k)}$ if and ony if $i \neq j$ and $\text{sim}(s_i^{(k)}, s_j^{(k)}) > \tau$ where sim represents some

similarity index between sentences (to be defined in the following paragraph), and $\tau \in [0, 1]$ is a similarity threshold,

- $w^{(k)} : E^{(k)} \to \mathbb{R}_+$ assigns edge weights $w_{ij}^{(k)} = \text{sim}(s_i^{(k)}, s_j^{(k)})$.

2.1.1 Similarity via Sentence Embeddings with SBERT

To capture sentence-level semantic meaning, and thus define the aforementioned sim function, we utilize Sentence-BERT (SBERT), a modification of the original Google's BERT (see Devlin et al. (2019)) architecture specifically designed to produce high-quality embeddings for sentence similarity tasks (Reimers and Gurevych (2019)). While BERT (Bidirectional Encoder Representations from Transformers) achieves state-of-the-art performance in many NLP benchmarks, it is aknowledged that it is not directly optimized for measuring sentence-to-sentence similarity. Its limitations are associated to that BERT produces contextualized token embeddings, not sentence embeddings. Aggregating these into a fixed-size representation (e.g., using mean pooling) yields inconsistent semantic distances. Furthermore, in similarity tasks, vanilla BERT requires joint input of any pair of two sentences which can be computationally inefficient for large-scale or many-to-many comparisons.

The SBERT modification addressed these issues by introducing a Siamese network architecture (see Chicco (2021)) that encodes each sentence independently using a shared transformer model. Sentence embeddings are then obtained by applying a pooling operation (typically mean-pooling) over the output token embeddings. Each sentence $s^{(k)} \in \mathcal{S}^{(k)}$ is mapped to a vector representation $\mathbf{v}_s^{(k)} \in \mathbb{R}^d$. The resulting fixed-size vectors lie in a semantically meaningful space, allowing cosine similarity to approximate semantic sentence similarity.

SBERT is fine-tuned on large-scale natural language inference (NLI) and semantic textual similarity (STS) datasets. This training encourages the model to place semantically similar sentences near each other in the embedding vector space. The resulting architecture supports fast, meaningful, and interpretable vector-based comparisons, making it highly suitable for applications in coherence evaluation, clustering, and semantic graph construction.

Given then an arbitrary pair of SBERT sentence embedding vectors $\mathbf{v}_{s_i^{(k)}}, \mathbf{v}_{s_j^{(k)}}$, the edge weight between them is defined as:

$$\mathbf{w}_{ij}^{(k)} = \text{sim}(s_i^{(k)}, s_j^{(k)}) := \frac{\langle \mathbf{v}_{s_i^{(k)}}, \mathbf{v}_{s_j^{(k)}} \rangle}{\|\mathbf{v}_{s_i^{(k)}}\|_2 \cdot \|\mathbf{v}_{s_j^{(k)}}\|_2}, \tag{1}$$

where $\langle \cdot, \cdot \rangle$ and $\| \cdot \|_2$ respectively denote the Euclidean inner product and norm. The embeddings are then used to define the Edge set $E^{(k)}$ and the final weight function $W^{(k)}$ that complete the construction of the TextRank-type graph; namely, the edge $e_{ij}^{(k)}$ is retained in $E^{(k)}$ if and only if $\mathbf{w}_{ij}^{(k)} > \tau$ for a fixed threshold $\tau \in [0.3, 0.7]$; optionally the top-k most similar neighbors could be retained for each node in order to enforce k-sparsity. The reduced weight function is then $w_{ij}^{(k)} := \begin{cases} \mathbf{w}_{ij}^{(k)}, \text{ if } \mathbf{w}_{ij}^{(k)} > \tau \\ 0, \text{ if } \mathbf{w}_{ij}^{(k)} \leq \tau \end{cases}$.

For each text d_k text in the corpus at hand, the resulting SBERT-based sentence based TextRank graph can be interpreted as a sentence-level semantic network. Its topological

and geometrical properties could reflect semantic coherence which could be recovered using graph curvature and/or entropy measures.

2.2 Ollivier–Ricci Curvature and Dominance on Sentence Graphs

The TextRank graph constructed over sentence embeddings with a geometric notion of curvature to reveal the structural properties of the sentence network for each text in the corpus. In particular, the notion of Ollivier-Ricci curvature Ollivier (2009) is employed, adapted to weighted graphs.

2.2.1 Metric Measure Space Structure

For $G^{(k)} = (V^{(k)}, E^{(k)}, w^{(k)})$ be the undirected weighted graph defined in the previous paragraph. To each node (sentence) $i \in V^{(k)}$ a discrete probability distribution $\mu_i^{(k)}$ supported its neighbors is defined via:

$$\mu_i^{(k)}(j) := \begin{cases} \frac{w_{ij}^{(k)}}{\sum_{m \in \mathcal{N}(i)} w_{im}^{(k)}}, & \text{for } j \in \mathcal{N}(i) \\ 0, & \text{for } j \notin \mathcal{N}(i) \end{cases},$$

where $\mathcal{N}(i)$ denotes the neighborhood of i inside $G^{(k)}$. This makes $(V^{(k)}, d_{G_{(k)}}, \{\mu_i^{(k)}, i \in V^{(k)}\})$ a discrete metric measure space, where $d_{G_{(k)}}(i,j)$ is the shortest-path distance on $G^{(k)}$, and $\{\mu_i^{(k)}, i \in V^{(k)}\}$ is the process on $V^{(k)}$ comprised of the abovementioned discrete measures.

2.2.2 Between Nodes Ollivier-Ricci Curvature

Following Ollivier (2009), the Ollivier–Ricci curvature $\kappa(i,j)$ along the edge (i,j) in $G^{(k)}$ is defined via optimal transport between the measures $\mu_i^{(k)}$ and $\mu_j^{(k)}$:

$$\kappa^{(k)}(i,j) := 1 - \frac{\mathcal{W}_1(\mu_i^{(k)}, \mu_j^{(k)})}{d_{G_{(k)}}(i,j)},$$

where $\mathcal{W}_1$ is the 1-Wasserstein distance between $\mu_i^{(k)}$ and $\mu_j^{(k)}$; the latter in the present discrete framework is defined as $\mathcal{W}_1(\mu_i^{(k)}, \mu_j^{(k)}) := \min_{\gamma \in \Gamma(\mu_i^{(k)}, \mu_j^{(k)})} \sum_{(m,n) \in \mathcal{N}_i \times \mathcal{N}_j} \gamma(m,n) d_{G_{(k)}}(m,n)$, where $\Gamma(\mu_i^{(k)}, \mu_j^{(k)})$ denotes the set of Borel probability distributions on $\mathcal{N}_i \times \mathcal{N}_j$ that have respective marginals $\mu_i^{(k)}, \mu_j^{(k)}$ (see Gao et al. (2017)). The computation of the Wasserstein distance and thus of between nodes curvature has a convenient LP representation-see for example Ch. 8 of Galichon (2016).

The curvature quantifies how well the neighborhoods of i and j are aligned. Positive curvature could be interpreted as local cohesion, while negative curvature may imply structural bottlenecks or bridges. Edges with negative curvature identify transitions across semantically distant or structurally weakly connected regions. By averaging $\kappa(i,j)$ over incident

edges, we define a sentence-level curvature score:

$$\bar{\kappa}^k(i) := \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \kappa^k(i,j),$$

which captures the local semantic structural integrity around sentence $s_i^{(k)}$. High positive values may indicate redundant or predictable contexts, while negative scores may indicate pivotal sentences or "non-smooth" transitions.

2.2.3 Step-Function Embedding of Curvature Profiles

Each sentence-level Ricci curvature vector $\bar{\kappa}^{(k)} = (\bar{\kappa}_1^{(k)}, \ldots, \bar{\kappa}_{n_k}^{(k)})$ is embedded into a piecewise-constant function defined on the unit interval:

$$\kappa_k(x) := \bar{\kappa}_i^{(k)} \quad \text{for } x \in \left(\frac{i-1}{n_k}, \frac{i}{n_k} \right], \quad i = 1, \ldots, n_k.$$

The resulting mapping-a standardized curvature profile-$\kappa_k : [0,1] \to \mathbb{R}$ is a right-continuous step function that preserves the position, magnitude, and sign of each curvature value in the sentence sequence. This embedding allows comparison between texts of varying length within a common function space, while capturing curvature dynamics over normalized discourse position.

2.2.4 Envelope Construction and Dominance Relation

Given the set of curvature profiles $\{\kappa_1, \ldots, \kappa_N\}$, define the corpus-wide envelope functions over $[0,1]$:

$$\overline{\kappa}(x) := \sup_k \kappa_k(x), \quad \underline{\kappa}(x) := \inf_k \kappa_k(x),$$

which represent, respectively, the pointwise best and worst structural cohesion observed at each normalized sentence position $x \in [0,1]$.

Furthermore, we say that text d_k dominates d'_k in structural cohesion, denoted $d_k \succcurlyeq d'_k$, if and only if:

$$\kappa_k(x) \geq \kappa_{k'}(x) \quad \text{for all } x \in [0,1].$$

This defines a pointwise dominance relation on the space of curvature profiles. It an interpretable (pre-) order: a dominating curvature profile never dips below in terms of curvature the dominated one across the normalized discourse. The $\overline{\kappa}$ curvature profile corresponds to an utopian (across the current corpus) curvature profile that may not be feasible by any text in the corpus. Dually, the $\underline{\kappa}$ corresponds to an dystopian minimal element of the corpus that may not correspond to any text there. Notice that the order may be incomplete; dominance requires the satisfaction of a potentially demanding functional inequality. Both envelopes however, as the following paragraph suggests, can be used in order to construct a text meta-ordering that essentially quantifies the obstacles to dominance.

2.3 Weighted Utopia Index (wUI)

To quantify how close a document's curvature profile comes to the corpus-wide best-case cohesion, we define the Weighted Utopia Index (wUI); this is an idea inspired from the Utopia Index considered in a quite different framework of income distributions and w.r.t. stochastic dominance relations (see Hadar and Russell (1969)) by Anderson et al. (2020). Given a $w_\theta : [0,1] \to \mathbb{R}_+$ that assumes the role of a nonnegative weighting function, The weighted shortfall of text d_k from the utopian upper envelope is given by:

$$\text{Shortfall}(d_k) := \int_0^1 w_\theta(x) \cdot [\overline{\kappa}(x) - \kappa_k(x)]_+ \, dx,$$

where $[\cdot]_+$ denotes the positive part. The maximal possible weighted shortfall is:

$$\text{MaxShortfall} := \int_0^1 w_\theta(x) \cdot [\overline{\kappa}(x) - \underline{\kappa}(x)] \, dx.$$

The Weighted Utopia Index is then defined as:

$$\text{wUI}(d_k) := 1 - \frac{\text{Shortfall}(d_k)}{\text{MaxShortfall}} = 1 - \frac{\int_0^1 w_\theta(x)[\overline{\kappa}(x) - \kappa_k(x)]_+ dx}{\int_0^1 w_\theta(x)[\overline{\kappa}(x) - \underline{\kappa}(x)] dx}.$$

This index lies in $[0,1]$, with:

- $\text{wUI}(d_k) = 1$ if $\kappa_k(x) = \overline{\kappa}(x)$ for all x, i.e., maximal structural cohesion.
- $\text{wUI}(d_k) = 0$ if $\kappa_k(x) = \underline{\kappa}(x)$ for all x, i.e., worst-case alignment.

The wUI thus defines an interpretable, and potentially human-informed (via the weight selection) cohesion measure over texts of arbitrary length, preserving curvature sign information. Mathematically, as the shortfall quantifies the obstacles that text d_k faces in order to establish utopian dominance across the corpus, wUI can be thought of as a meta-ordering that functions as a completion to $\succcurlyeq$; d_k is preferred according to wUI from d'_k, iff it achieves a greater wUI-the ordering is complete. A sufficient condition for that is $d_k \succcurlyeq d_j$.

By construction, the wUI evaluated at an arbitrary document in the corpus contains information about the curvature profiles that inhabit the corpus, through the envelope profiles. Changing the set of texts may alter the wUI score of a fixed text due to the potential alterations of the envelopes; this suggests the potential sensitivity of the index on the inclusion in the corpus of texts with "extreme" curvature profiles.

2.3.1 Training the Weight Function via Ordered Probit

The weighting function remains as a fudge factor in the construction of wUI above. In this paper we use an exponential-quadratic parametric form, $w_\theta(x) = \exp(\theta_0 + \theta_1 x + \theta_2 x^2)$ where the parameters $\theta = (\theta_0, \theta_1, \theta_2) \in \mathbb{R}^3$ are estimated from human coherence ratings in the COHESENTIA dataset-Maimon and Tsarfaty (2023), via an ordered probit model that links latent utility to observed ordinal preferences. This exponential-quadratic form is chosen over,

say, a Beta density or spline-based function given that it is smooth, unimodal, and closed under log-concavity—a desirable inductive bias for capturing primacy or closure emphasis.

Specifically, let the dataset contain K documents (speeches, essays, or passages), each rated by a set of annotators $a = 1, \ldots, A$. Each annotation is an ordinal label $r_k^{(a)} \in \{1, 2, \ldots, L\}$, representing increasing perceived coherence. We posit a latent variable model of the form:

$$y_k^{(a)*} := \mathrm{wUI}_k(\theta) + \varepsilon_k^{(a)},$$

where $y_k^{(a)*}$ is the unobserved coherence utility of document k for annotator a, $\varepsilon_k^{(a)} \sim \mathcal{N}(0, 1)$ is standard Gaussian noise, and $\mathrm{wUI}_k(\theta)$ is the model-determined weighted Utopia Index for document d_k.

The observed label $r_k^{(a)}$ is assumed to arise via discretization of this latent utility:

$$r_k^{(a)} = l \quad \Longleftrightarrow \quad \tau_{l-1} < y_k^{(a)*} \leq \tau_l, \quad l = 1, \ldots, L,$$

where $\{\tau_l\}_{l=1}^{L-1}$ are monotonic cutpoints with $\tau_0 := -\infty$ and $\tau_L := +\infty$. The probability of label l is:

$$\mathbb{P}[r_k^{(a)} = l \mid \theta] = \Phi(\tau_l - \mathrm{wUI}_k(\theta)) - \Phi(\tau_{l-1} - \mathrm{wUI}_k(\theta)),$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function.

The log-likelihood across all annotations is:

$$\mathcal{L}(\theta, \{\tau_l\}) = \sum_{a=1}^{A} \sum_{k=1}^{K} \log \left[\Phi(\tau_{r_k^{(a)}} - \mathrm{wUI}_k(\theta)) - \Phi(\tau_{r_k^{(a)}-1} - \mathrm{wUI}_k(\theta)) \right].$$

Maximum likelihood estimation is performed over θ. We adopt a fixed-threshold ordinal probit model by uniformly spacing C - 1 cutpoints on the standardized latent utility scale (from -1.5 to 1.5), which avoids identifiability issues and simplifies estimation by focusing optimization exclusively on the weight parameters governing the Utopia Index. Optimization can be carried out using quasi-Newton methods or constrained trust-region solvers.

This estimation procedure infuses the weighting function with empirical information about where human raters place value in the rhetorical development of a text. If, for instance, coherence early in a document is systematically rewarded, the learned weight function will place greater mass on early $x \in [0, 1]$, thus encouraging curvature profiles that are structurally consistent at the beginning of the discourse.

The estimated weight function $w_{\theta^*}(x)$ provides insight into how human annotators implicitly value cohesion over rhetorical time; $\theta_1 < 0$ may indicate strong preference for early cohesion—consistent with anchoring, topic scaffolding, and the primacy effect. $\theta_1 > 0$ may suggest emphasis on late structure—possibly aligned with persuasive closure or narrative arcs. Nonzero θ_2 could be aligned with U-shaped or bell-shaped preferences, indicating mid-speech importance.

2.4 Average Adjacent SBERT Coherence

In our pipeline, and as a complement to the WUI, we employ the SBERT-based average weight across adjacent nodes of our TextRank-type graph;

$$\texttt{SBERT-Coherence}(d_k) = \frac{1}{n_k - 1} \sum_{i=1}^{n-1} w_{i\,i+1}^{(k)}.$$

This score captures the extent to which adjacent sentences in the speech share similar semantic content, reflecting smooth rhetorical progression. High scores indicate fluent, well-aligned transitions; lower scores suggest abrupt shifts or topic discontinuities. It leverages SBERT's pretrained semantic knowledge, hopefully providing a sequential/adjacent complement to curvature-based global cohesion scores.

2.5 A Nash-type Multiobjective Cohesion Score

Scalarizing the resulting vector score (wUI, `SBERT-Coherence`), define a final cohesion score using a Nash bargaining-type product as:

$$\mathcal{N}(d_k) := \mathrm{wUI}(d_k)^{w_1} \cdot \texttt{SBERT-Coherence}(d_k)^{w_2}, \tag{2}$$

where $w = (w_1, w_2) \in \Delta^2$ is a probability vector representing trade-offs between the two constitutional cohesion scores.

The Nash-type Cohesion Score can be justified via a multiobjective ordering of the texts inside the corpus based on the pair of the aforementioned scores; d_k is said to pairwise dominate d'_k, iff $\mathrm{wUI}(d_k) \geq \mathrm{wUI}(d'_k)$ and $\texttt{SBERT-Coherence}(d_k) \geq \texttt{SBERT-Coherence}(d'_k)$. Then the text that maximizes the Nash Cohesion Score can be easily seen to be a Pareto efficient element of the corpus according to this ordering; no gain in wUI or Coherence-BERT is possible without a loss in the other. Furthermore, the Nash product is invariant under positive affine transformations of each criterion, preserving comparability across scales.

This formulation can be particularly suited for comparative rhetorical analysis, where differing strengths in local vs global cohesion may suggest distinct communicative strategies or ideological styles.

2.6 Cohesion Analysis Summary Algorithm

The following pseudocode summarizes our full pipeline for computing cohesion metrics from a corpus of texts. It integrates graph-theoretic structure extraction, curvature analysis, sentence-level semantic alignment, a trained version of the Utopia Index using human-labeled data, and a Nash-type multiobjective scoring.

Algorithm 1: Text Cohesion Analysis via Curvature–Weighted UI–SBERT–Nash Pipeline

Input: Speech corpus $\mathcal{D} = \{d_1, \ldots, d_N\}$ with sentence segmentation
Output: Cohesion metrics per speech: curvature, SBERT-coherence, UI, Nash

foreach *speech* $d \in \mathcal{D}$ **do**
- Extract ordered sentences $\mathcal{S} = \{s_1, \ldots, s_n\}$;
- Compute SBERT embeddings $\mathbf{v}_i \leftarrow \texttt{SBERT}(s_i)$;
- Construct similarity graph $G = (V, E, w)$ with $w_{ij} = \cos(\mathbf{v}_i, \mathbf{v}_j)$;
- Prune edges below similarity threshold τ;
- Compute Ollivier–Ricci curvature $\kappa(i, j)$ via optimal transport;
- Derive node-level curvature scores $\bar{\kappa}_i = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \kappa(i, j)$;
- Construct step-function curvature profile $\kappa_k(x)$ over $x \in [0, 1]$;

Construct upper and lower envelope profiles:

$$\overline{\kappa}(x) = \sup_k \kappa_k(x), \quad \underline{\kappa}(x) = \inf_k \kappa_k(x)$$

Train exponential weight function $w_\theta(x) = \exp(\theta_0 + \theta_1 x + \theta_2 x^2)$ via ordinal probit likelihood:

$$y_k^* = \mathrm{wUI}_k(\theta) + \varepsilon_k, \quad y_k = j \iff \tau_{j-1} < y_k^* \leq \tau_j$$

foreach *speech* d_k **do**
- Compute weighted Utopia Index:

$$\mathrm{wUI}_k = 1 - \frac{\int_0^1 w_\theta(x) \cdot [\overline{\kappa}(x) - \kappa_k(x)]_+ \, dx}{\int_0^1 w_\theta(x) \cdot [\overline{\kappa}(x) - \underline{\kappa}(x)] \, dx}$$

- Compute SBERT coherence score:

$$\texttt{SBERT-Coherence}(d_k) = \frac{1}{n_k - 1} \sum_{i=1}^{n_k - 1} \cos\left(\mathbf{v}_i, \mathbf{v}_{i+1}\right)$$

- Compute Nash Cohesion Score:

$$\mathcal{N}(d_k) = \sqrt{\mathrm{wUI}(d_k) \cdot \texttt{SBERT-Coherence}(d_k)}$$

Aggregate metrics across $\mathcal{D}$ and rank speeches accordingly;

Implementation Details. All computations were implemented in Python. Sentence embeddings were generated using `sentence-transformers` (`all-MiniLM-L6-v2`). SBERT coherence scores were computed as the mean cosine similarity between adjacent sentence embeddings, directly capturing semantic alignment in rhetorical progression.

Semantic graphs were constructed using pairwise cosine similarities, with edges pruned below a similarity threshold $\tau = 0.3$. Ollivier–Ricci curvature was computed using the Wasserstein Distance via the `POT` library, yielding a curvature series over sentence positions.

3 Benchmarking on the Cohesentia Dataset

To evaluate the effectiveness of our proposed Weighted Utopia Index (UI) in modeling rhetorical cohesion, we benchmark it against human-annotated coherence scores using the Cohe-

sentia dataset (see Maimon and Tsarfaty (2023)). This section briefly introduces the dataset, and presents evaluation results that align machine-derived cohesion scores with human judgments—highlighting discourse-theoretic implications and motivating our use of top-k evaluation metrics.

3.1 The Cohesentia Dataset

The Cohesentia corpus (refer to Maimon and Tsarfaty (2023)) is a curated benchmark made up of 500 generated narratives. On average, each narrative contains 6.5 sentences, and each has undergone a coherence evaluation by at least two MTurk raters. Each document is segmented into sentences and labeled on a 5-point ordinal scale:

$$y_k \in \{1 : \text{very incoherent},\ 2,\ 3,\ 4,\ 5 : \text{very coherent}\},$$

reflecting perceived logical progression, structural consistency, and rhetorical fluency, regardless of topic or style. Annotations include raw text and sentence segmentation, as well as human-assigned holistic and incremental coherence scores.

As an illustration, consider story ID `5`, a short reflective text about London. The dataset entry includes the following metadata:

- **Title:** *What is your favorite place, city or country of the world? Why?*
- **Text:** "My favorite place to live, work and play is of course London! It has a lot to offer both the city dweller and tourist alike. The atmosphere in the cosmopolitan metropolis feels like home. Additionally, I love the abundance of culture here; from its vibrant nightlife to its research-friendly universities. As for my favorite country/city there's no definitive answer – it just depends on what I'm feeling that day or week!"
- **Sentence count:** 5
- **Holistic coherence score:** 3.0
- **Incremental coherence score:** 5.0

The example demonstrates a moderately cohesive narrative structure with strong incremental development, as reflected in the human based average score divergence (holistic: 3.0, incremental: 5.0).

Cohesentia is the first public dataset—introduced in this work—to combine sentence-level structural features with fine-grained human annotations on cohesion. It enables a direct evaluation of algorithmic cohesion measures over structurally varied texts.

3.2 Learned Weight Function for Weighted UI

The parametric exponential-square weight function described in the methodology is estimated in the training part of the dataset. The analysis focuses on the human holistic coherence score. The ML estimated parameters $\hat{\theta}$ via the ordered probit model are:

$\hat{\theta} = (-1.27 \times 10^{-4}, \quad 4.65, \quad -5.0)$, implying a strongly front-loaded weight function: $w^*_{\theta}(x) = \exp(-0.00013 + 4.65x - 5x^2)$.

This curve peaks early-to-mid in textual time, emphasizing initial structural cohesion and decays sharply-consistent with cognitive and rhetorical salience of earl to mid coherence. The learned emphasis on early to mid coherence is partly consistent with the primacy effect which posits that initial information disproportionately shapes interpretive judgments (Glanzer and Cunitz (1966)).

3.3 Quantitative Evaluation Against Human Judgments

We benchmark the three scores on the Cohesentia dataset: SBERT Coherence Score, Weighted Utopia Index (wUI), and Nash Cohesion Score. Performance is reported over four metrics, including top-20 restricted variants of rank correlation. The following table presents the results showing in blue, best performance for any given criterion:

Table 1: Benchmarking Cohesion Metrics Against Human Annotations (Cohesentia)

Model	ρ (full)	τ (full)	ρ (top-20)	τ (top-20)	Top-20 Agreement	nDCG
SBERT Coherence Score	0.061	0.042	-0.060	-0.050	35%	0.899
Weighted Utopia Index	0.008	0.006	**0.139**	**0.117**	**45%**	**0.911**
Nash Cohesion Score	**0.062**	**0.047**	0.099	0.083	40%	0.897

Regarding the evaluation metrics we use full-sample Spearman ρ and Kendall τ correlations that repesent ordinal agreement across all speeches. Top-20 correlations isolate model alignment within the most coherent texts as judged by annotators. Top-20 Agreement measures overlap in the top 20 model vs. human rankings. Normalized Discounted Cumulative Gain (nDCG) rewards correct ordering of high-quality items, emphasizing ranked retrieval performance-see Wang et al. (2013).

In order to handle annotation uncertainty in mid-rank regions, we follow practices in evaluation literature-see for example Järvelin and Kekäläinen (2002), and report both full and top-20 metrics. Top-20 coherence rankings can be especially informative, as agreement across annotators seems highest for saliently coherent texts.

The Weighted UI outperforms all baselines on top-20 correlations, agreement, and nDCG. Its negative full-rank correlations may reflect noise in the middle of the label distribution, but its superior top-20 precision indicates its potential alignment with salient coherence judgments. The Nash cohesion score does not uniformly outperform all competing metrics, yet it demonstrates a consistent stabilizing effect. Specifically, it achieves the highest full-sample correlations (Spearman ρ and Kendall τ) and reliably improves upon the weaker-performing metric in each evaluation subset. For instance, it corrects the top-20 ranking failures of SBERT coherence (which exhibits negative ρ and τ) while approximating the stronger performance of the weighted Utopia Index. This behavior is consistent with the interpretation of Nash aggregation as a multiobjective compromise: given the exponents equality, it balances distinct coherence criteria.

The computational pipeline regarding the full (train and test) set of the 500 (small) texts in the dataset was performed in approximately 2 minutes on a laptop with the Apple M1

chip and 8Gb of memory.

4 Case Study: Cohesion Across Inaugural Speeches

To illustrate the explanatory value of our cohesion metrics, we restrict attention to six post-1950 English-language US-presidents inaugural addresses. This subset controls for language, translation variability, and major rhetorical-historical shifts, allowing clearer comparison. The weight function enabled in the construction of wUI, and subsequently of the Nash scores, is the one trained in the Cohesentia dataset. The primary texts were sourced from publicly accessible archives, including the Avalon Project at Yale Law School and official governmental websites such as the United States White House (`whitehouse.gov`). The results are illustrated in the following table:

Table 2: Speech Cohesion Summary: Best 2 = blue, Worst = red.

Speech	**SBERT-CS**	**wUI**	**NASH**
Obama (2009)	0.3275	0.4889	0.4001
Kennedy (1961)	0.3745	0.6281	0.4850
Clinton (1993)	0.3300	0.5711	0.4341
Reagan (1981)	0.3193	0.2777	0.2977
Bush Jr. (2005)	0.3509	0.7742	0.5212
Trump (2025)	0.3318	0.5924	0.4433

The post-1950 English-language inaugural speeches reveal notable variation in cohesion strategies across leaders. Bush Jr.'s 2005 address emerges as the strongest overall, achieving the highest Weighted Utopia Index (UI) and Nash Score, suggesting a speech with structurally coherent curvature profiles that align well with early rhetorical emphasis, as learned from the Cohesentia dataset. Kennedy's 1961 speech ranks first regarding the SBERT-CS and second across across the remaining metrics. Clinton, Obama and Trump occupy a middle ground, with Trump exhibiting particularly high structural cohesion and Clinton showing well-balanced performance across both dimensions. Obama's 2009 speech ranks lower in both semantic and structural terms, potentially due to a fluid rhetorical progression that may deprioritize early/mid cohesion. Reagan's 1981 address receives the lowest scores, reflecting limited alignment with both structural and semantic coherence. Obviously, the current wUI and Nash results depend crucially on the particular corpus of speeches, as well as on the early-to mid text emphasis placed by the trained weight function at hand.

5 Discussion and Future Directions

This paper introduced a framework for analyzing text cohesion, leveraging tools from discrete geometry, curvature embeddings dominance, and statistical learning. At its core, the methodology constructs a sentence-level semantic graph based on SBERT embeddings and cosine similarity, over which Ollivier-Ricci curvature is computed to quantify local rhetorical

flow and structural alignment. The resulting curvature profiles are interpreted as normalized time series reflecting how semantic connectivity evolves throughout a speech.

Building on this geometric foundation, we propose a new numerical score—the Weighted Utopia Index (wUI)—to assess each speech's alignment with corpus-level structural norms. The wUI leverages a dominance logic: it compares the speech's curvature profile to a reference envelope constructed from lower and upper empirical bounds across the corpus. Crucially, the wUI is computed via a learned weighting function, which adjusts temporal sensitivity by assigning greater or lesser importance to different rhetorical regions.

The weight parameters are estimated by minimizing the negative log-likelihood of an ordinal probit model trained on the Cohesentia dataset—an annotated corpus of speeches labeled with human coherence ratings. This yields an interpretable utility-based model where the latent perceived coherence is a function of the speech's weighted UI. The probit structure reflects both ordinal judgments and underlying cognitive models of coherence formation. To our knowledge, this is the first application of trained, dominance-based weighting schemes to textual structure analysis in NLP—extending the theoretical framework of Anderson et al. (2020) beyond economic distributions into the semantic domain.

The final metric, the Nash Cohesion Score, synthesizes geometric alignment with sentence-level semantic flow, serving as a Pareto compromise.

Methodological Extensions. Several natural extensions arise from this framework. First, while the weighting function $w_\theta(x)$ is currently specified as an exponential-quadratic form, it could be generalized using nonparametric or spline-based formulations to capture more nuanced patterns of rhetorical emphasis. Second, the ordinal probit model could be expanded into hierarchical or latent trait structures to account for annotator heterogeneity or domain-specific effects. Finally, training on larger and more diverse corpora—beyond Cohesentia—could improve generalizability, enabling fine-tuned models for political, legal, or pedagogical discourse.

Taken together, the combined use of discrete (network) curvature, dominance theory, statistical training, and machine learning could offer a potentially extensible framework for the numerical analysis of textual cohesion. Beyond individual speech analysis, the proposed methodology is well-suited to larger comparative studies. For instance, applying the framework to annual State of the Union addresses may reveal structural shifts in rhetorical cohesion across administrations, political alignments, or in response to geopolitical crises. Similar applications can be envisioned for campaign speeches, UN General Assembly addresses, or inaugural declarations.

Finally, future extensions could aim to normalize cohesion measures with respect to the inherent complexity of the discussed subject matter. This may involve incorporating Minimum Description Length (MDL) principles-see Rissanen (1987)-as proxy for conceptual complexity. A potential complexity-normalized cohesion index could evaluate how well cohesion is maintained despite increasing semantic or informational complexity. Such an index would be especially relevant in more technical domains—such as central bank monetary policy reports, regulatory briefings, or economic outlooks—where linguistic clarity and cohesion must coexist with content density.

Data Availability

For the Cohesentia corpus refer to Maimon and Tsarfaty (2023). The texts of the inaugural speeches were sourced from the publicly accessible archives: a. Avalon Project at Yale Law School (`avalon.law.yale.edu`), and b. the United States White House (`whitehouse.gov`).

Conflicts of Interest

There are no conflicts of interest.

Funding

No funding was received for this article.

References

Amancio, D. R., O. N. Oliveira Jr, and L. d. F. Costa (2012): "Using complex networks to quantify consistency in the use of words," *Journal of Statistical Mechanics: Theory and Experiment*, 2012, P03005.

Anderson, G., T. Post, and Y.-J. Whang (2020): "Somewhere between utopia and dystopia: Choosing from multiple incomparable prospects," *Journal of Business & Economic Statistics*, 38, 502–515.

Barzilay, R. and M. Lapata (2008): "Modeling local coherence: An entity-based approach," *Computational Linguistics*, 34, 1–34.

Barzilay, R. and L. Lee (2004): "Catching the drift: Probabilistic content models, with applications to generation and summarization," *NAACL*, 4, 113–120.

Burstein, J., K. Kukich, S. Wolff, C. Lu, L. Braden-Harder, and M. Harris (1998): "Automated scoring using a hybrid feature identification technique," in *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics,*, 206–210.

Chicco, D. (2021): "Siamese neural networks: An overview," *Artificial neural networks*, 73–94.

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2019): "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.

Galichon, A. (2016): *Optimal transport methods in economics*, Princeton University Press.

Gao, R., X. Chen, and A. Kleywegt (2017): "Distributional robustness and regularization in statistical learning," *arXiv preprint*.

Glanzer, M. and A. R. Cunitz (1966): "Two storage mechanisms in free recall," *Journal of verbal learning and verbal behavior*, 5, 351–360.

Grosz, B. J., A. K. Joshi, and S. Weinstein (1995): "Centering: A framework for modeling the local coherence of discourse," *Computational linguistics*, 21, 203–225.

Hadar, J. and W. R. Russell (1969): "Rules for ordering uncertain prospects," *The American economic review*, 59, 25–34.

Jamaati, M. and A. Mehri (2018): "Text mining by Tsallis entropy," *Physica A: Statistical Mechanics and its Applications*, 490, 1368–1376.

Järvelin, K. and J. Kekäläinen (2002): "Cumulated gain-based evaluation of IR techniques," *ACM Transactions on Information Systems (TOIS)*, 20, 422–446.

Kwon, H., H. T. Kwon, and W. C. Yoon (2016): "An information-theoretic evaluation of narrative complexity for interactive writing support," *Expert Systems with Applications*, 53, 219–230.

Lin, Z., H. T. Ng, and M.-Y. Kan (2011): "Automatically evaluating text coherence using discourse relations," in *Proceedings of the 49th Annual Meeting of the Association for Computational Lin- guistics: Human Language Technologies. Associa- tion for Computational Linguistics*, 997–1006.

Louis, A. and A. Nenkova (2012): "A coherence model based on syntactic patterns," in *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, 1157–1168.

Maimon, A. and R. Tsarfaty (2023): "Cohesentia: A novel benchmark of incremental versus holistic assessment of coherence in generated texts," *arXiv preprint arXiv:2310.16329*.

Mann, W. C. and S. A. Thompson (1988): "Rhetorical structure theory: Toward a functional theory of text organization," in *Text—Interdisciplinary Journal for the Study of Discourse*, De Gruyter Mouton, vol. 8, 243–281.

Mihalcea, R. and P. Tarau (2004): "Textrank: Bringing order into text," in *Proceedings of the 2004 conference on empirical methods in natural language processing*, 404–411.

Ollivier, Y. (2007): "Ricci curvature of metric spaces," *Comptes Rendus Mathematique*, 345, 643–646.

——— (2009): "Ricci curvature of Markov chains on metric spaces," *Journal of Functional Analysis*, 256, 810–864.

Papadimitriou, C., K. Karamanos, F. Diakonos, V. Constantoudis, and H. Papageorgiou (2010): "Entropy analysis of natural language written texts," *Physica A: Statistical Mechanics and its Applications*, 389, 3260–3266.

Reimers, N. and I. Gurevych (2019): "Sentence-bert: Sentence embeddings using siamese bert-networks," *arXiv preprint arXiv:1908.10084*.

Rissanen, J. (1987): "Stochastic complexity and the MDL principle," *Econometric Reviews*, 6, 85–102.

Soricut, R. and D. Marcu (2006): "Discourse generation using utility-trained coherence models," in *Proceedings of the COLING/ACL*, 803–810.

Wang, Y., L. Wang, Y. Li, D. He, and T.-Y. Liu (2013): "A theoretical analysis of NDCG type ranking measures," in *Conference on learning theory*, PMLR, 25–54.

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS
ΣΧΟΛΗ ΟΙΚΟΝΟΜΙΚΩΝ ΕΠΙΣΤΗΜΩΝ SCHOOL OF ECONOMIC SCIENCES
ΤΜΗΜΑ ΟΙΚΟΝΟΜΙΚΗΣ ΕΠΙΣΤΗΜΗΣ DEPARTMENT OF ECONOMICS

Department of Economics
Athens University of Economics and Business

List of Recent Working Papers

2023

01-23 Real interest rate and monetary policy in the post Bretton Woods United States, George C. Bitros and Mara Vidali

02-23 Debt targets and fiscal consolidation in a two-country HANK model: the case of Euro Area, Xiaoshan Chen, Spyridon Lazarakis and Petros Varthalitis

03-23 Central bank digital currencies: Foundational issues and prospects looking forward, George C. Bitros and Anastasios G. Malliaris

04-23 The State and the Economy of Modern Greece. Key Drivers from 1821 to the Present, George Alogoskoufis

05-23 Sparse spanning portfolios and under-diversification with second-order stochastic dominance, Stelios Arvanitis, Olivier Scaillet, Nikolas Topaloglou

06-23 What makes for survival? Key characteristics of Greek incubated early-stage startup(per)s during the Crisis: a multivariate and machine learning approach, Ioannis Besis, Ioanna Sapfo Pepelasis and Spiros Paraskevas

07-23 The Twin Deficits, Monetary Instability and Debt Crises in the History of Modern Greece, George Alogoskoufis

08-23 Dealing with endogenous regressors using copulas; on the problem of near multicollinearity, Dimitris Christopoulos, Dimitris Smyrnakis and Elias Tzavalis

09-23 A machine learning approach to construct quarterly data on intangible investment for Eurozone, Angelos Alexopoulos and Petros Varthalitis

10-23 Asymmetries in Post-War Monetary Arrangements in Europe: From Bretton Woods to the Euro Area, George Alogoskoufis, Konstantinos Gravas and Laurent Jacque

11-23 Unanticipated Inflation, Unemployment Persistence and the New Keynesian Phillips Curve, George Alogoskoufis and Stelios Giannoulakis

12-23 Threshold Endogeneity in Threshold VARs: An Application to Monetary State Dependence, Dimitris Christopoulos, Peter McAdam and Elias Tzavalis

13-23 A DSGE Model for the European Unemployment Persistence, Konstantinos Giakas

14-23 Binary public decisions with a status quo: undominated mechanisms without coercion, Efthymios Athanasiou and Giacomo Valletta

15-23 Does Agents' learning explain deviations in the Euro Area between the Core and the Periphery? George Economides, Konstantinos Mavrigiannakis and Vanghelis Vassilatos

16-23 Mild Explocivity, Persistent Homology and Cryptocurrencies' Bubbles: An Empirical Exercise, Stelios Arvanitis and Michalis Detsis

17-23 A network and machine learning approach to detect Value Added Tax fraud, Angelos Alexopoulos, Petros Dellaportas, Stanley Gyoshev, Christos Kotsogiannis, Sofia C. Olhede, Trifon Pavkov

18-23 Time Varying Three Pass Regression Filter, Yiannis Dendramis, George Kapetanios, Massimiliano Marcellino

19-23 From debt arithmetic to fiscal sustainability and fiscal rules: Taking stock, George Economides, Natasha Miouli and Apostolis Philippopoulos
20-23 Stochastic Arbitrage Opportunities: Set Estimation and Statistical Testing, Stelios Arvanitis and Thierry Post
21-23 Behavioral Personae, Stochastic Dominance, and the Cryptocurrency Market, Stelios Arvanitis, Nikolas Topaloglou, and Georgios Tsomidis
22-23 Block Empirical Likelihood Inference for Stochastic Bounding: Large Deviations Asymptotics Under *m*-Dependence, Stelios Arvanitis and Nikolas Topaloglou
23-23 A Consolidation of the Neoclassical Macroeconomic Competitive General Equilibrium Theory via Keynesianism (Part 1 and Part 2), Angelos Angelopoulos
24-23 Limit Theory for Martingale Transforms with Heavy-Tailed Noise, Stelios Arvanitis and Alexandros Louka

2024

01-24 Market Timing & Predictive Complexity, Stelios Arvanitis, Foteini Kyriazi, Dimitrios Thomakos
02-24 Multi-Objective Frequentistic Model Averaging with an Application to Economic Growth, Stelios Arvanitis, Mehmet Pinar, Thanasis Stengos, Nikolas Topaloglou
03-24 State dependent fiscal multipliers in a Small Open Economy, Xiaoshan Chen, Jilei Huang, Petros Varthalitis
04-24 Public debt consolidation: Aggregate and distributional implications in a small open economy of the Euro Area, Eleftherios-Theodoros Roumpanis
05-24 Intangible investment during the Global Financial Crisis in the EU, Vassiliki Dimakopoulou, Stelios Sakkas and Petros Varthalitis
06-24 Time will tell! Towards the construction of instantaneous indicators of different agent-types, Iordanis Kalaitzoglou, Stelios Arvanitis
07-24 Norm Constrained Empirical Portfolio Optimization with Stochastic Dominance: Robust Optimization Non-Asymptotics, Stelios Arvanitis
08-24 Asymptotics of a QLR-type test for optimal predictive ability, Stelios Arvanitis
09-24 Strongly Equitable General Equilibrium Allocations, Angelos Angelopoulos
10-24 The Greek macroeconomy: A note on the current situation and future outlook, Apostolis Philippopoulos
11-24 Evolution of Greek Tax System, A Survey of Legislated Tax Changes from 1974 to 2018, Panagiotis Asimakopoulos
12-24 Macroeconomic Impact of Tax Changes, The case of Greece from 1974 to 2018, Panagiotis Asimakopoulos
13-24 `Pareto, Edgeworth, Walras, Shapley' Equivalence in a Small Economy, Angelos Angelopoulos
14-24 Stimulating long-term growth and welfare in the U.S, James Malley and Apostolis Philippopoulos
15-24 Distributionally Conservative Stochastic Dominance via Subsampling, Stelios Arvanitis
16-24 Before and After the Political Transition of 1974. Institutions, Politics, and the Economy of Post-War Greece, George Alogoskoufis
17-24 Gaussian Stochastic Volatility, Misspecified Volatility Filters and Indirect Inference Estimation, Stelios Arvanitis, Antonis Demos
18-24 Endogenous Realtor Intermediation and Housing Market Liquidity, Miroslav Gabrovski, Ioannis Kospentaris, Victor Ortego-Marti
19-24 Universal Choice Spaces and Expected Utility: A Banach-type Functorial Fixed Point, Stelios Arvanitis, Pantelis Argyropoulos, Spyros Vassilakis

2025

01-25 Democracy, redistribution, and economic growth: Some evidence from post-1974 Greece, George C. Bitros

02-25 Assessing Downside Public Debt Risks in an Environment of Negative Interest Rates Growth Differentials, Yiannis Dendramis, Georgios Dimitrakopoulos and Elias Tzavalis

03-25 Wicksellian General Equilibrium, Angelos Angelopoulos

04-25 Institutions and Ethics, Angelos Angelopoulos

05-25 General Equilibrium in a Capitalist Economy, Angelos Angelopoulos

06-25 Public investment multipliers revisited: The role of production complementarities, Vasiliki Dimakopoulou, George Economides, and Apostolis Philippopoulos

07-25 The heterogeneous causal effects of the EU's Cohesion Fund, Angelos Alexopoulos, Ilias Kostarakos, Christos Mylonakis, and Petros Varthalitis

08-25 Gaussian Invariant Markov Chain Monte Carlo, Michalis K. Titsias, Angelos Alexopoulos , Siran Liu, and Petros Dellaportas

09-25 Financial Intermediation, Investment, and Monetary Policy, Ioannis Kospentaris and Florian Madison

10-25 Temperature as a premium factor. Insurance incorporating climate change. A comparative case study between Europe and North America, Nomikou-Lazarou Irene

11-25 How High will our Buildings Become? Angelos Angelopoulos

12-25 Emergency Room General Equilibrium, Angelos Angelopoulos

13-25 Updating Density Estimates using Conditional Information Projection: Stock Index Returns and Stochastic Dominance Relations, Stelios Arvanitis, Richard McGee, and Thierry Post

14-25 Measuring Textual Cohesion via Graph Curvature, Weighted Dominance, and Embedding-Based Coherence, Stelios Arvanitis

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

ΣΧΟΛΗ ΟΙΚΟΝΟΜΙΚΩΝ ΕΠΙΣΤΗΜΩΝ
SCHOOL OF ECONOMIC SCIENCES

ΤΜΗΜΑ ΟΙΚΟΝΟΜΙΚΗΣ ΕΠΙΣΤΗΜΗΣ
DEPARTMENT OF ECONOMICS

Department of Economics
Athens University of Economics and Business

The Department is the oldest Department of Economics in Greece with a pioneering role in organising postgraduate studies in Economics since 1978. Its priority has always been to bring together highly qualified academics and top quality students. Faculty members specialize in a wide range of topics in economics, with teaching and research experience in world-class universities and publications in top academic journals.

The Department constantly strives to maintain its high level of research and teaching standards. It covers a wide range of economic studies in micro-and macroeconomic analysis, banking and finance, public and monetary economics, international and rural economics, labour economics, industrial organization and strategy, economics of the environment and natural resources, economic history and relevant quantitative tools of mathematics, statistics and econometrics.

Its undergraduate program attracts high quality students who, after successful completion of their studies, have excellent prospects for employment in the private and public sector, including areas such as business, banking, finance and advisory services. Also, graduates of the program have solid foundations in economics and related tools and are regularly admitted to top graduate programs internationally. Three specializations are offered:1. Economic Theory and Policy, 2. Business Economics and Finance and 3. International and European Economics. The postgraduate programs of the Department (M.Sc and Ph.D) are highly regarded and attract a large number of quality candidates every year.

For more information:

https://www.dept.aueb.gr/en/econ/