A Spectral Analysis of Connectivity in Trade Networks

Miles Guerin

Advised by:
Professor Yannis M. Ioannides
Professor James M. Murphy

A Senior Honors Thesis
Tufts University
May 2025

Abstract

This thesis examines the evolution of trade connectivity and country centrality from 1950 to 2014. Using bilateral trade data from Fouquin and Hugot (2016), we construct yearly weighted directed graphs of global trade flows and apply a range of graph-theoretic methods to quantify changes in network structure over time. We focus on network connectivity, defined as the degree to which countries are mutually reachable through trade relationships and the extent to which the network resists fragmentation. Specifically, we use spectral methods, such as the directed Laplacian and approximations of graph conductance, to quantify changes in network connectivity. We provide a comparison of these metrics and report their values for both our observed trade networks and for null models of the network that preserve each country's out-degree and distribution of edge weights. Using these methods, we do not identify major shifts in the connectivity of aggregate flows in the world trade network. However, we do identify a sharp increase in connectivity within the intra-European subgraph during the early 1990s. Finally, we extend this methodology to select product-level and product category-level trade networks.

Keywords: International Trade Network, Spectral Graph Theory, Network Connectivity

JEL: F14, F15, N70

Acknowledgments

First and foremost, I would like to thank my advisors, Professor Yannis Ioannides and Professor James Murphy, for their invaluable guidance, mentorship, and support throughout this project. This thesis has also greatly benefited from discussions with Christopher Coscia, Lenore Cowen, Elena Elbarmi, Michael Klein, Kabir Tripathi, and Calvin Zhang. I am especially thankful to the attendees of the Network Science in Economics Conference for their feedback on this thesis. In particular, I would like to acknowledge Hiroyuki Asai, Alex Hayes, Matthew Jackson, Karl Rohe, James Stratton, and Yajie Xu for their comments and suggestions. I also extend my gratitude to the Office of Scholar Development for financially supporting my attendance at the conference. Finally, I would like to thank my family and fiancée for their unwavering support and encouragement.

1 Introduction

Globalization has fundamentally reshaped international trade, creating a vast, interconnected network of economic relationships among nations. This complex web of trade interactions cannot be fully understood through traditional economic analyses, which often focus on trade volumes, bilateral balances, or isolated economic indicators (Fagiolo, 2010). Such approaches miss the influence of interdependencies and the overall structure of international trade on outcomes. To address these limitations, a subset of the international trade literature has applied network theory to examine the structure of international trade networks, to explore the centrality of certain countries in trade, the formation of trade communities, and the diffusion of shocks across the trade network (Chaney, 2014; De Benedictis et al., 2014; Fagiolo, 2016; Hoang, Piccardi, & Tajoli, 2023).

This paper employs network theory to analyze two key aspects of the international trade network: centrality and modularity. Centrality metrics indicate the influence of individual countries within the trade network, reflecting how some countries function as major trade hubs, while others occupy more peripheral roles. Metrics like PageRank centrality, which accounts for both direct and indirect connections, allow us to identify countries that have a disproportionate impact on global trade flows due to their central position within the network (Deguchi et al., 2014; Hoang, Piccardi, & Tajoli, 2023).

Modularity, on the other hand, highlights the degree to which the trade network is partitioned into distinct clusters or communities. In international trade, high modularity may indicate the presence of regional trade blocs or tightly knit groups of countries with dense internal trade and fewer external connections, while low modularity reflects a more integrated global system. We choose to analyze the dynamics of modularity to better explain how geopolitical events, such as the dissolution of the USSR, impact the overall cohesiveness and regionalization of the trade network. Metrics of modularity act as a barometer of the globalization-regionalization balance in trade networks. High modularity suggests a fragmented network dominated by regional or local clusters, which may point to a trend of regionalization or trade protectionism. Conversely, low modularity indicates a more integrated global system, where countries trade freely between blocs. Furthermore, modularity can measure the resilience of the trade network to shocks. For instance, if a modular trade network is highly dependent on a single cluster or hub, a disruption in that cluster—such as economic sanctions, a natural disaster, or political instability—could ripple through the entire global system. In contrast, a less modular network may more easily absorb the shock to a single node or collection of trade relationships.

The core objective of this paper is to map the structural evolution of international trade from 1950 to 2014 by examining changes in the centrality of specific countries and the overall connectivity of the trade network. Using historical trade data from Fouquin and Hugot (2016), we construct directed weighted networks for each year of data, where each network represents aggregate flows of trade between countries measured in British Pounds. Given this series of temporal network snapshots, we attempt to measure the rise and fall of influential countries, the formation and dissolution of trade blocs, and the effects of major geopolitical events on the structure of international trade. To this end, we utilize key methodologies frequently used in the computer science literature: PageRank centrality and spectral graph theory in order to discuss changes in network centrality and connectivity over time.

PageRank centrality enables us to capture the influence of countries beyond raw trade volume. In international trade, PageRank centrality has been utilized by Deguchi et al. (2014); Kim and Yun (2022); and Hoang, Piccardi, and Tajoli, (2023) across different time periods and levels of data to analyze the dynamics of countries becoming more or less central in the network of international trade.

In contrast, spectral graph theory has been less commonly applied in economics but is widely utilized in other disciplines, including computer science for tasks like clustering and image segmentation, physics for studying complex systems, and biology for analyzing protein interaction networks and gene regulatory systems. In the study of international trade, Laenen and Sun (2020) employed spectral graph theory to cluster countries into four trading clusters, but otherwise, little work has been done on determining trade modularity through spectral graph theory. That being said, clustering and the formation of trade communities have been explored quite extensively in the international trade literature, namely: Fagiolo (2010); Chaney (2014); and Abbate et al. (2018). We chose to use spectral graph theory in this analysis because, in addition to its utility for community detection, spectral methods provide a means to determine the algebraic connectivity of a graph (Bapat, 2010). In the context of international trade, a higher algebraic connectivity indicates a well-integrated trade network where countries are strongly interconnected, making it less susceptible to fragmentation. Conversely, lower algebraic connectivity suggests that the trade network is more modular or loosely connected, potentially reflecting the presence of distinct trade blocs or regional clusters.

To summarize, this paper offers a longitudinal analysis of the international trade network by focusing on the evolution of centrality and modularity from 1950 to 2014. We examine how shifts in global economic power, regional trade dynamics, and major geopolitical events have reshaped the structure of international trade. Then, concerning the methods used, we provided a detailed breakdown of each of their constructions, a comparison between them, and a discussion of under what settings each metric can be applied. Lastly, we benchmark observed trade connectivity against basic null models of international trade and briefly extend the analysis to networks product- and commodity-level trade.

2 Theoretical Background

In this section, we explore the theoretical underpinnings of international trade literature by reviewing two prominent frameworks: gravity models and network theory. These models provide complementary perspectives on the mechanisms that drive trade flows and the structural characteristics of international trade networks.

To model international trade, gravity models are commonly used to explain the volume of trade between countries. The gravity model, first introduced by Tingergen (1962), borrows from Newton's law of gravitation, stating that the trade flow between two countries is directly proportional to their economic sizes (measured by GDP) and inversely proportional to the geographic distance between them. The naive gravity model is formulated as:

$$X_{ij} = G \cdot Y_i^{\beta_1} Y_j^{\beta_2} D_{ij}^{-\beta_3} \cdot \varepsilon_{ij} \ ,$$

where X_{ij} is the aggregate bilateral trade from country i to country j, Y_i and Y_j are the GDP of country i and j respectively, D_{ij} is the distance between country i and j under some metric, and ϵ_{ij} is a normally distributed error term with a mean of 0. Numerous augmentations have been made to this model to account for variables that can further influence bilateral trade between countries, such as the inclusion of common language, common borders, and joint membership in trade agreements (Kabir et al., 2017). For empirical applications, the main strength of gravity models lie in their ability to econometrically estimate the effect of trade agreements, common currencies, and other factors on the bilateral flow of trade (Baier et al., 2014). In practice, these models have been quite successful in explaining the relationship between country characteristics and variations in bilateral trade flows. However, gravity models are limited in their ability to capture the network effects present in international trade, where trade flows between countries are not independent but rather influenced by the global structure of trade relationships. This is where network theory provides a valuable complement that enables us to consider the centrality of certain countries and the overall interdependence of trade.

Network theory offers an approach to analyzing international trade that emphasizes the interdependence of trade flows and the broader structure of the international trade network. In this framework, countries are represented as nodes, and yearly trade flows are represented by linkages from exporter to importer countries. Furthermore, weights may be assigned to these linkages, or edges, to represent heterogeneous flows of trade between countries. Repeating such a construction year by year results in a time series of networks. Formally, for each year t, the set of nodes is the set of countries, $N_t = \{1, ..., n_t\}$, and define the edge weight, $w_{ij}^t \geq 0$ to be the value of exports from country i to country j in year t where $i, j \in N_t$. Note that due to the asymmetry of imports and exports, w_{ij}^t is not necessarily equal to w_{ij}^t. Given this set of countries, N_t and a matrix representing all the directed edge weights, W_t, we can represent the network of aggregate trade in year t as a weighted directed graph.

If we wanted to consider trade flows by commodity, bilateral flows could be broken down into industry-specific flows, $w^t_{ij}(h)$ where $h \in \{1, ..., H\}$ indicates a specific commodity (Fagiolo, 2016). This decomposition results in H different networks for each year, t, and the time series of all these networks is called a multilayer network (Kivelä et al., 2014). Alternatively, if we wanted to model flows between M industries across different countries, each node in the network would be represented as a country-industry pair. Let $\mathcal{I} = \{1, ..., M\}$ be the set of all industries, then the flow from industry $k \in \mathcal{I}$ in country i to industry $l \in \mathcal{I}$ in country j in year t is given by $w^t_{(i,k)\to(j,l)}$ or $w^t_{ik,jl}$.

Regardless of the type of network constructed, an assortment of metrics is commonly used to analyze the topological properties of trade networks and their dynamics. For example, graph connectivity measures the extent to which countries are linked within the trade network, either directly or indirectly. A highly connected network indicates that goods can flow efficiently between countries, reducing barriers and enhancing resilience against localized disruptions. In contrast, a network with low connectivity might suggest potential bottlenecks or dependencies on specific trade routes or partners, making the network more vulnerable to disruptions. Beyond the overall connectivity of the network, we may also be concerned with the centrality of individual countries within the network. Different types of centrality—such as degree centrality, betweenness centrality, and eigenvector centrality—can measure various types of influence on the overall trade graph. Degree centrality, for instance, measures the number of direct trade connections a country has, indicating how integrated it is within the network. Betweenness centrality, on the other hand, measures how often a country lies on the shortest path between pairs of countries, potentially representing a country pivotal in facilitating indirect trade. Eigenvector centrality considers not only the number of connections but also the influence of those connections, highlighting countries that are well-connected to other influential trade partners. Beyond centrality and connectivity, we are particularly interested in the clustering of certain communities in the network and the overall modularity of the trade network.

Modularity, or connectivity, in network theory refers to the degree to which a network can be partitioned into distinct communities or clusters. In the context of international trade, these clusters may represent regional trade blocs or groups of countries with dense internal trade links but relatively sparse external connections. High modularity indicates a fragmented network with distinct trade communities, while low modularity suggests a more integrated global trade system. Despite significant trade globalization, the world trade network remains a largely modular network as trade becomes clustered around large advanced economies (Ioannides, 1996; Fagiolo, 2016).

2.1 Random Graphs and Null Models

Beyond the construction of trade networks from data, the mechanisms behind their generation have been explored using tools from random graph theory. First defined by Erdös and Renyi (1959), random graphs are used to study the properties of networks where the edges between nodes are determined probabilistically. This modeling approach is suited for systems where the network structure plays a critical role in shaping outcomes but cannot be fully predicted from the initial parameters. In particular, random graph theory has been used extensively by mathematicians, sociologists, computer scientists, and biologists in analyzing social and biological networks. In the domain of economics, Ioannides (2015) has provided an extensive overview of various types of random graphs and their respective usages in economics; and has provided a model for random graphs of trade in differentiated products between countries. While most models of random graphs assume the number of nodes to be static, Jackson (2011) reviews the construction of growing random networks in which nodes are added to the graph dynamically and form edges with existing nodes. Such a model of random graphs extends beyond the analysis of this paper, but provides a means to simulate the introduction of new countries or firms to the network of international trade over time. In particular, Setayesh, et al. (2022) have utilized exponential random graph models to explore the structural, economic, geographical, and political factors that influence the evolution of the trade network.

Beyond modeling, random graphs can be used to assess the statistical significance of a metric or structural feature of the actual graph (Jackson, 2008; De Clerck et al., 2024). In the context of international trade, comparing the observed trade network to a simulated distribution of random graphs enables a form of hypothesis testing. This approach allows us to determine whether a change in the trade network, such as shifts in centrality or graph modularity, can be attributed to an external factor, such as a political or economic shock, or if it is indistinguishable from random fluctuations (noise) within the network. By comparing the observed metric to its distribution in randomized versions of the network, we can determine whether the observed change is statistically significant or likely a result of random chance. In constructing random graphs, or null models, Fagiolo (2016) emphasizes the importance of carefully selecting which features of the world trade network are preserved during the randomization process. The constraints on these features should strike a balance: they must be strict enough to retain a sense of realism while remaining flexible enough to generate a meaningful distribution of the metric of interest for hypothesis testing.

3 Data

We use data from Fouquin and Hugot (2016), which compiles the British Pound value of total nominal exports between pairs of countries between 1950 and 2014. For this time period, the data is primarily sourced from IMF data and supplemented with primary sources from trading countries. The trade amounts are at the country-pair year level. We drop non-sovereign ISO-3 codes from the dataset. For each year, we have between 93 and 196 countries with recorded trade flows. Across all years, we have over 1.8 million bilateral trade observations. The total volumes of all exports for each year are shown in Figure 22. Given the dramatic rise in international trade over the previous decades, it is important to move beyond mere volumetric analysis and instead examine the network structure of trade.

One difficulty of analyzing trade data—particularly for earlier time periods—is the occurrence of missing data, as not all countries report their trade flow. Sajedianfard et al. (2021) explore the impact of the missing data on metrics of network centrality. The authors conclude that since network behavior in the trade network is motivated by leading and major participating countries that consistently report their trade flow; missing data from smaller non-reporting countries has no significant impact on the structure of the network. This means we can draw inferences from measures of network centrality and modularity, despite the occurrence of missing data.

In our analysis of trade in products and product categories, we will be using data from Gaulier and Zignago (2025), in which they have compiled bilateral trade flows for more than 200 countries and 5000 products as categorized by Harmonized System Codes (HS codes) for the years 1995-2023. Flows are given in both quantities and dollar amounts for each year. Products are aggregated using the first six digits of the HS codes. For our analysis, we will be further aggregating to either the first two, or the first four digits of HS codes, depending on the product or product category we are analyzing.

3.1 Construction of the Trade Network

For each year between 1950 and 2014, we construct a network where nodes represent countries, and edges represent exports. Let $N_t = \{1, ..., n_t\}$ be the set of n_t countries that we have trading data for in year t. Now, for each year, define the weight matrix, W_t, as an $n_t \times n_t$ matrix with matrix entries denoted as w_{ij}^t, where w_{ij}^t represents the total value of exports (in British Pounds) from country i to country j in year t, where $i, j \in N_t$. $w_{ij}^t = 0$ implies there is no value of recorded exports from country i to j in the given year, and $w_{ij}^t > 0$ implies some value of exports was recorded from country i to country j in year t. As we are considering aggregate flows between countries, and not intra-country inter-industry flows, we set $w_{ii}^t = 0$ for all countries.

As network summary statistics, we plot the number of countries in the network, the number of non-zero edges, and edge density (as computed by the number of non-zero edges divided by $(N_t * (N_t - 1))$ for each year between 1950 and 2014 are shown in Figure 23. These summary statistics are chosen to replicate Bhattacharya's (2007) summary of international trade networks using a different dataset. The number of countries represented by the trade graph increased from 93 in 1950 to 185 in 2014, capturing the creation of new states and more existing states reporting their export totals. The number of non-zero edges and edge density also increased since 1950, indicating both more countries reporting trade flows and countries trading with a wider array of other countries.

The weight matrices for the years 1950 and 2014 are displayed as logarithmic heatmaps in Figure 24 and Figure 25, respectively. The rows and columns are ordered by the total exports of each country, so that the largest exporters occupy the first positions. Since the value of the exports of country i to country j is, by construction, equal to the value of the imports of country j from country i, the transpose of the export weight matrix W_t, denoted $(W_t)^\intercal$, corresponds to the import matrix, where entry (i, j) gives the value of the imports received by country i from country j in year t. The difference between country i's exports to country j and imports from country j is thus given by $w_{ij}^t - w_{ji}^t$ (Since w_{ji}^t is equal to the ij entry of $(W_t)^\intercal$). The matrix of trade imbalances, defined as $w_{ij}^t - w_{ji}^t$, represents the balance of trade for each pair of countries. Positive values indicate that country i exports more to country j than it imports from them, while negative values suggest the opposite. As an example, we plot this matrix for the world network of trade in 2014 in Figure 26. This matrix illustrates how some countries export more in the majority of their bilateral connections, whereas other countries import more in the majority of their bilateral trades. These imbalances in trade flows mean the set of countries that are central as importers can differ significantly from the set of countries that are central as exporters. Countries with consistent trade surpluses—where $w_{ij}^t > w_{ji}^t$ across many trading partners—exhibit systematically higher out-degree strengths in the export network. This directional dominance can increase a country's PageRank centrality, particularly when its export flows are to other well-connected nodes. As a result, persistent net exporters often emerge as structural hubs in the directed network, while net importers may be positioned more peripherally depending on the centrality metric used (Deguchi et al., 2014; Chaney, 2014).

3.2 Intra-European Trade Network

Due to potential issues of missing data in the complete trade network, we also consider a smaller trade network in which all exporting and importing countries are located on the European continent. This intra-European network is defined by geographic inclusion rather than political or institutional membership; that is, we include all countries geographically situated in Europe, regardless of their affiliation with organizations such as the European Union (EU) or its predecessor institutions like the European Economic Community (EEC). Analyzing this network in particular allows us to observe how changes in connectivity and centrality of the network coincide with major political and economic events that occurred on the European continent between 1950 and 2014 such as, the postwar reconstruction, the formation and expansion of the European Common Market, the fall of the Berlin Wall, the reunification of Germany, the collapse of the Soviet Union, and the eastward enlargement of the EU. Not only have all these events resulted in major political shifts, but they have each reshaped the economic geography of Europe. For example, the dissolution of the USSR and the resulting entry of many Eastern European countries into the European trade network in the early 1990s increased the number of trade linkages across the continent. Or consider how the establishment of the eurozone contributed to further trade integration in Western Europe (Baldwin, 2006). Given the existing work done on major structural shifts in intra-European trade, when analyzing the evolution of connectivity in this sub-network, we will have a good frame of reference for interpreting shifts in our metrics.

In order to construct the intra-European trade network, like above, for each year between 1950 and 2014, we construct directed weighted networks representing the value of exports from one European country to another. We plot the same summary statistics as above for the intra-European trade network in Figure 27. In the summary statistics, we observe a steady rise in non-zero edges for most of the period, and then a discrete jump in both the number of nodes and non-zero edges in the early 1990s following the dissolution of the USSR. The export weight matrices for 1950 and 2014 are displayed as logarithmic heatmaps in Figure 28 and Figure 29, respectively. The difference between these matrices shows the expansion of intra-European trade from 1950 to 2014, as, by 2014, nearly all possible bilateral trade connections between countries on the continent are populated. Visualizations of the intra-European trade network for 1950 and 2014 are shown in Figure 1 and Figure 2 respectively.

Network of Intra-European Trade in 1950 (Two Largest Export Destinations).

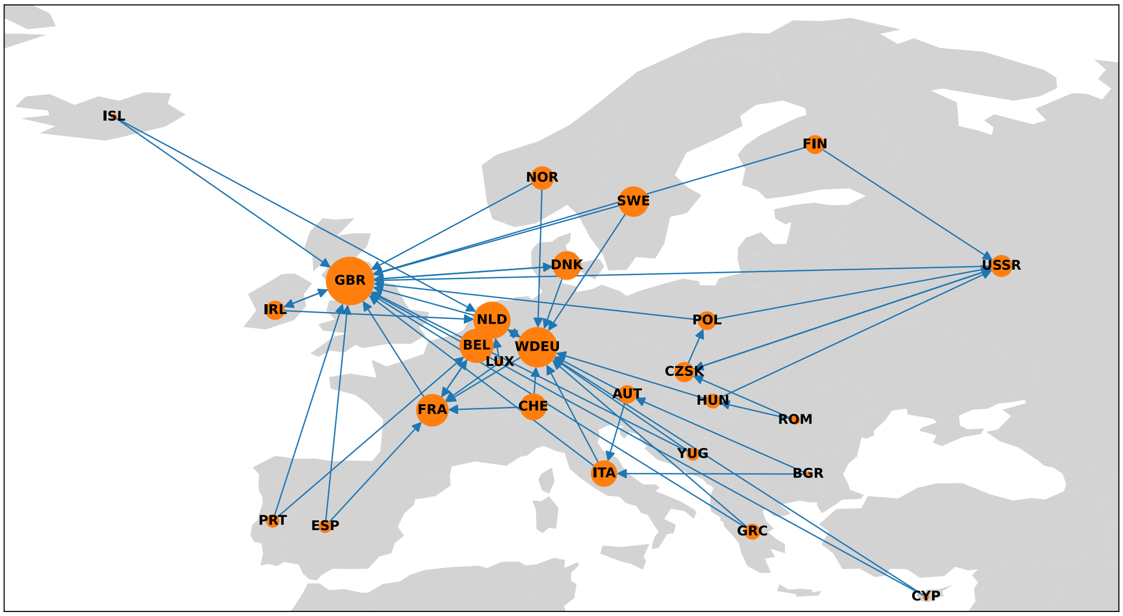

Figure 1: Using the 1950 network of intra-European trade, directed edges are drawn for each country's two largest export destinations. Country sizes are based on their in-degree, representing the value of goods they import.

Network of Intra-European Trade in 2014 (Two Largest Export Destinations).

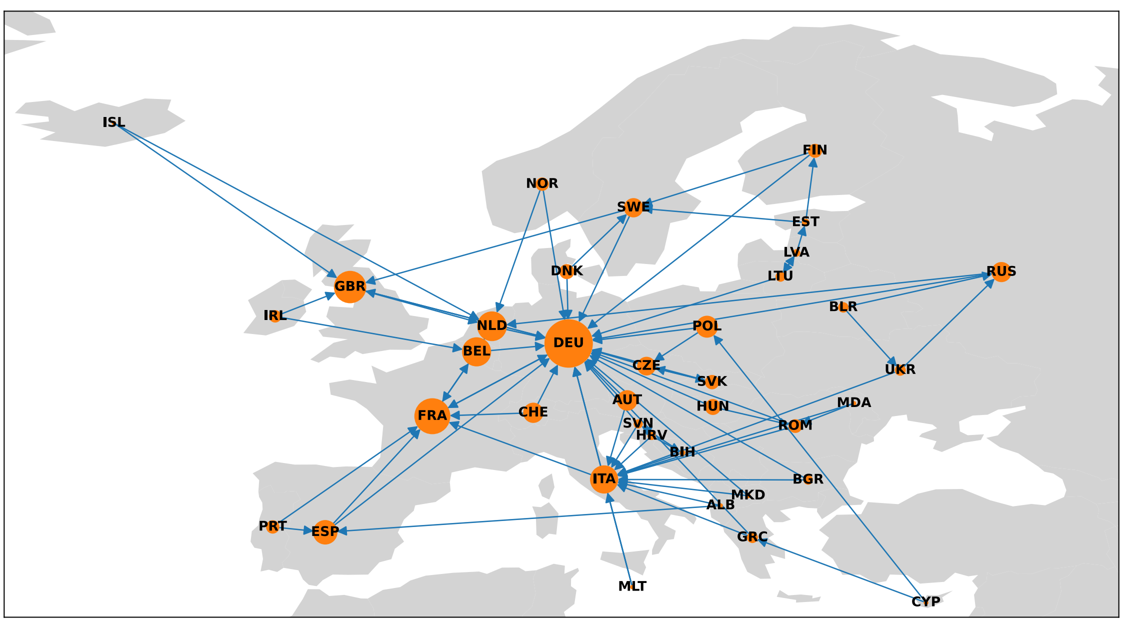

Figure 2: Using the 2014 network of intra-European trade, directed edges are drawn for each country's two largest export destinations. Country sizes are based on their in-degree, representing the value of goods they import.

4 Methodology and Results

In this section, we utilize various measures of graph centrality and modularity to infer how the structure of the trade network has evolved between 1950 and 2014. To analyze centrality, we use the PageRank algorithm developed by Brin and Page (1998). In our analysis of graph modularity, we will employ spectral methods involving the graph Laplacian and an approximation of graph conductance.

4.1 Centrality Analysis

PageRank evaluates the importance of a country not only by its direct connections but also by the significance of its connections' links. This means a country is deemed central not merely by the volume of its trade, but by how central its trade partners are within the network. Zhang et al. (2021) apply weighted PageRank to World Input-Output networks, and introduce a tuning parameter to adjust for the relative importance of weights and degree in the model. Hoang, Piccardi, and Tajoli (2023) utilize weighted and unweighted PageRank to analyze bilateral trade volumes between 1996 and 2019. The authors interpret a decrease in the network's average PageRank, among other metrics, as an indication of network homogenization, as large countries' PageRanks decrease relative to small economies.

To analyze the evolution of centrality in the networks we constructed above, we utilize weighted PageRank with the damping factor set to 0.95. Given the weight matrix of exports, W_t, we start by taking the transpose of the matrix such that w_{ij}^t now represents the values of the imports from country i to country j. We take the transpose of the matrix so that our computed PageRank scores will capture which countries are the most central as exporters, rather than importers. The PageRank value of country i in year t is denoted as $PR(i)_t$, and is defined recursively as:

$$PR(i)_t = \alpha \sum_{j \in N_t,\, j \neq i} w_{ij}^t \frac{PR(j)_t}{d_j^{\text{out}}} + \frac{1-\alpha}{|N_t|},$$

where,

- N_t is the set of countries in year t.

- $\alpha \in (0, 1]$ is the damping factor. While $\alpha \geq 0.85$ is standard in PageRank analysis (Brin and Page, 1998), we use $\alpha = 0.99$ to ensure the computed PageRank scores remain reflective of the true network structure.
- $d_j^{\text{out}} = \sum_{j \in N_t, j \neq i} w_{ij}^t$ is the weighted out-degree of country j.
- $PR(j)_t$ is the weighted PageRank of country j in year t.

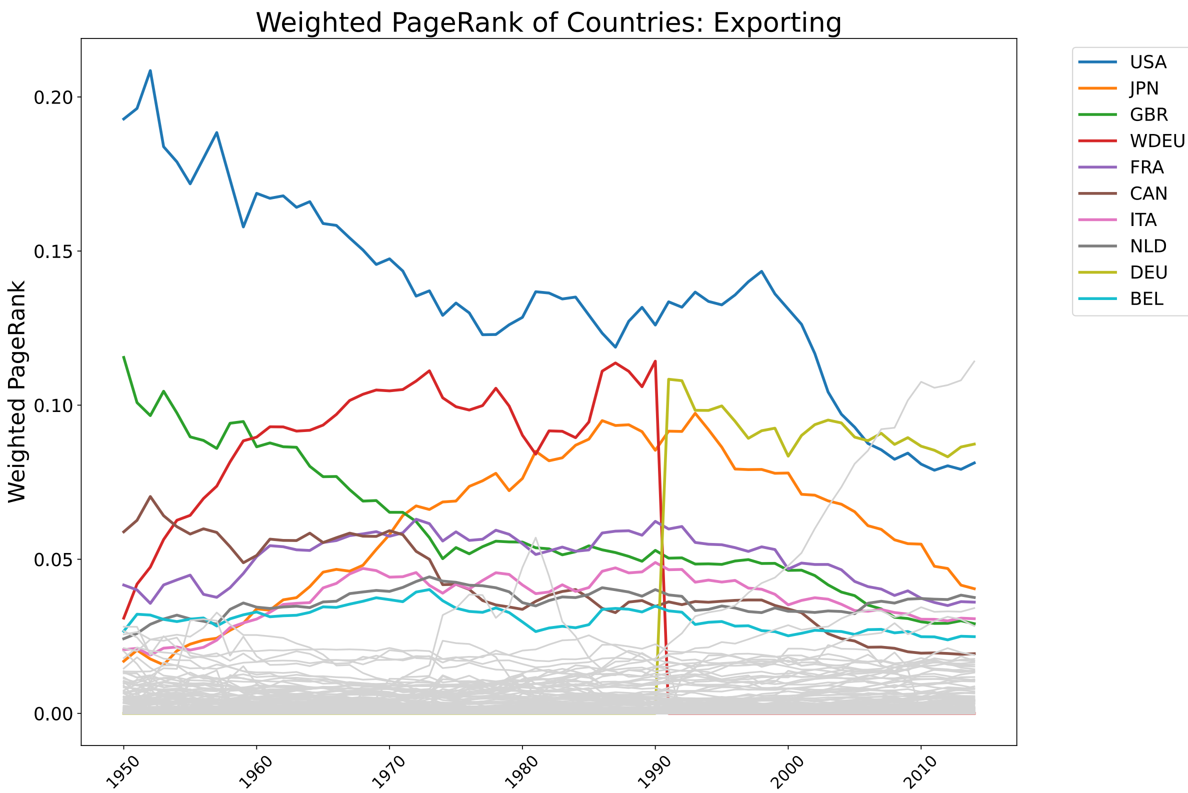

Figure 3: Weighted PageRank of countries in terms of exports with $\alpha = 0.99$. Countries with the top 10 average PageRank values are labeled.

The chart of each country's weighted PageRank (PR) in each year for $\alpha = 0.99$ is displayed in Figure 3. The increase in West Germany's (WDEU) PR in the 1950s and early 1960s reflects the recovery of the West German economy under the Marshall Plan, and the country's exports to GDP ratio more than doubled in the 1950s (Vonyó, 2018). Then, with the subsequent reunification of Germany, we observe a flattening of its weighted PR. A similar pattern emerges for Japan, which experienced significant economic growth in the post-WW2 period, but then experienced prolonged stagflation. Over the full period, we observe the decline in Great Britain's (GBR) PageRank with the collapse of the British Empire. The US also appears to have declined in export centrality significantly since the 1950s. While this is in part due to shifts in global manufacturing, this can also be attributed to a massive spike in the US's PR following the conclusion of WW2, which can be seen in the PRs of the full sample of data (1827-2014) which we include in Figure 30 in the appendix. In addition to plotting the PR scores for exports, we also plot the PR scores for countries in terms of import centrality in Figure 31, which we compute by not taking the transpose W and then calculating weighted PR scores. The charts are fairly similar for the top 10 countries, apart from the US, which maintains a similar PR throughout the sample.

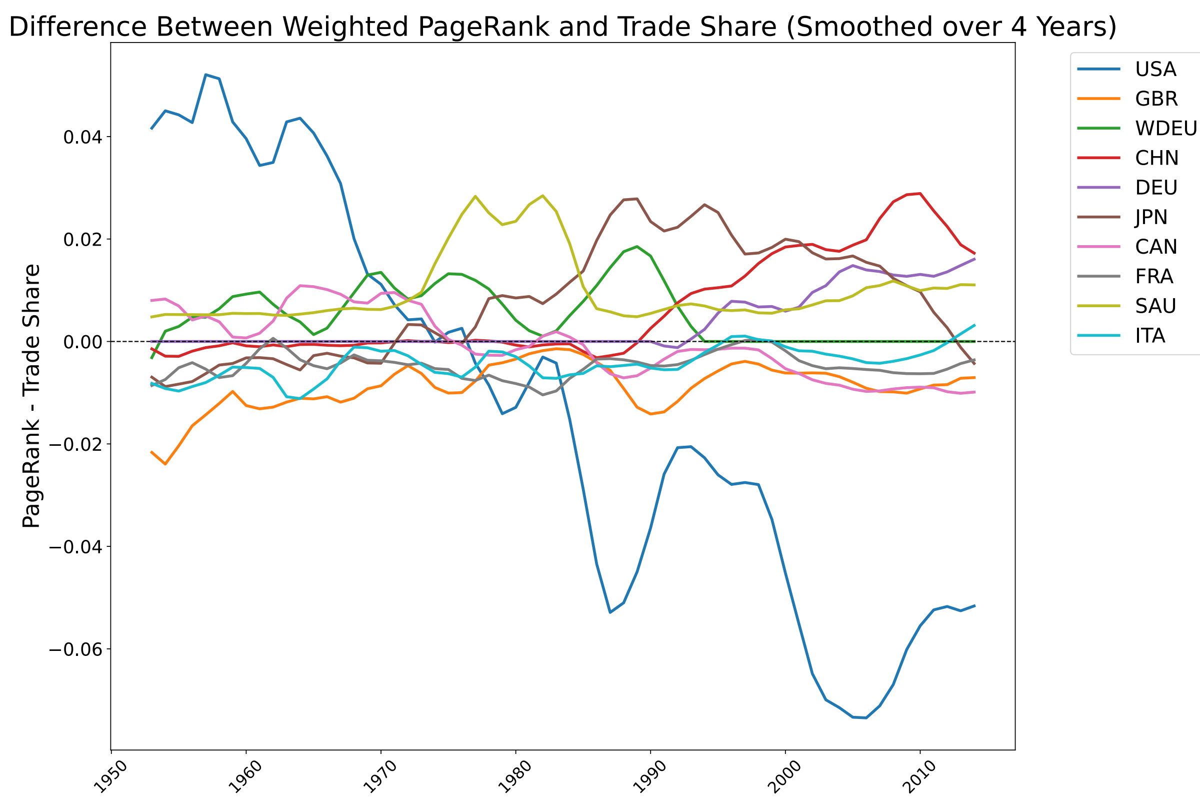

Figure 4: Weighted PageRank - Export Share for top 10 countries of the global trade network. PageRanks are computed using the transpose of the network's weight matrix. $\alpha = 0.99$

Naturally, we might ask whether or not countries' exporting PageRank scores differ significantly from their respective shares of exports in the network. PageRank is highly correlated with export shares since as exports increase, w_{ij} increases, which increases $PR(i)$. So, if we look at the difference between a country's weighted exporting PR and its share of total exports in the network, we can observe which countries are more central despite how little they export, and which countries are less central in the network despite how much they export. In Figure 4 for the ten countries with the highest average export PR and for each year, we plot the difference between each country's PR and their fraction of world exports given by,

$$\text{Country } i\text{'s Export Share} = \frac{\sum_{j \in N} w_{ij}}{\sum_{i \in N} \sum_{j \in N} w_{ij}} .$$

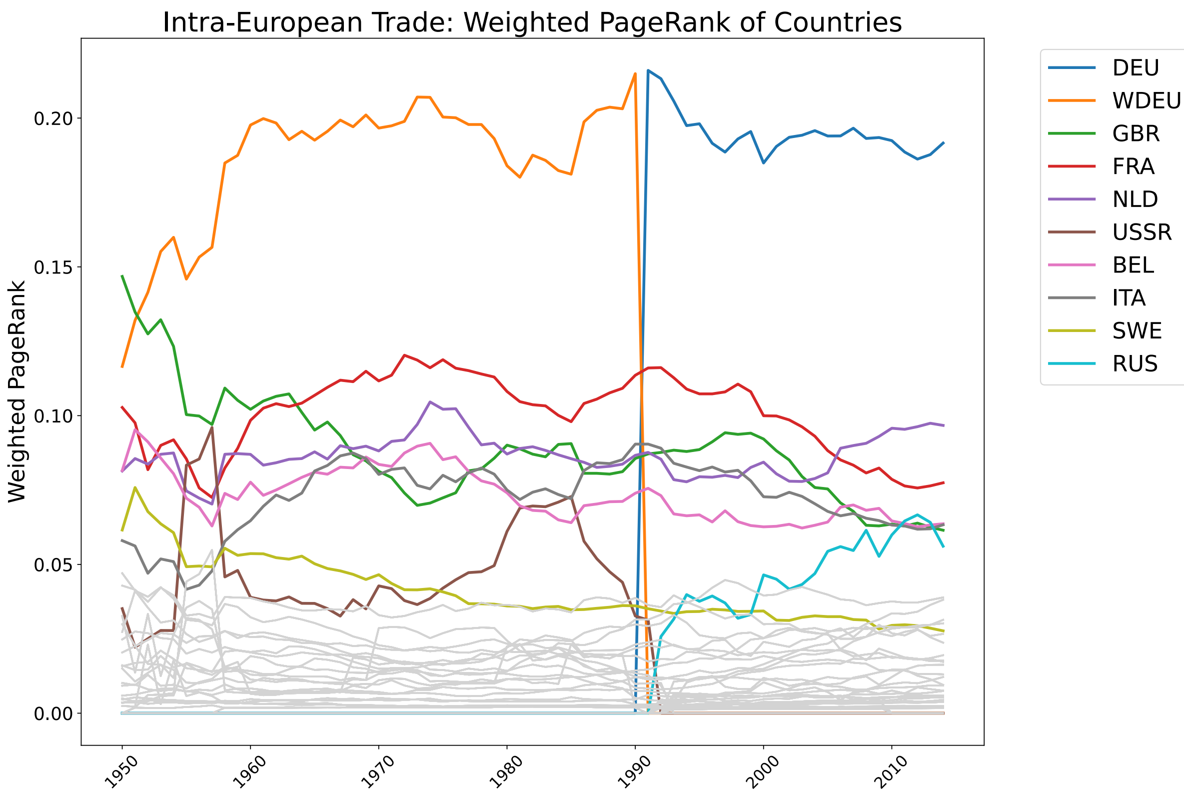

Figure 5: Weighted PageRank of countries in terms of exports for intra-European trade $\alpha = 0.99$. Countries with the top 10 average PageRank values are labeled.

Charting this difference over time reveals how a country's centrality within the global trade network can diverge from its share of total exports. For instance, the US initially had a PageRank centrality that exceeded its export share, reflecting its outsized influence in the postwar trade network. However, in the decades following, its centrality has declined relative to its share of global exports, suggesting that its position within the export network has become less central despite continued high export volumes.

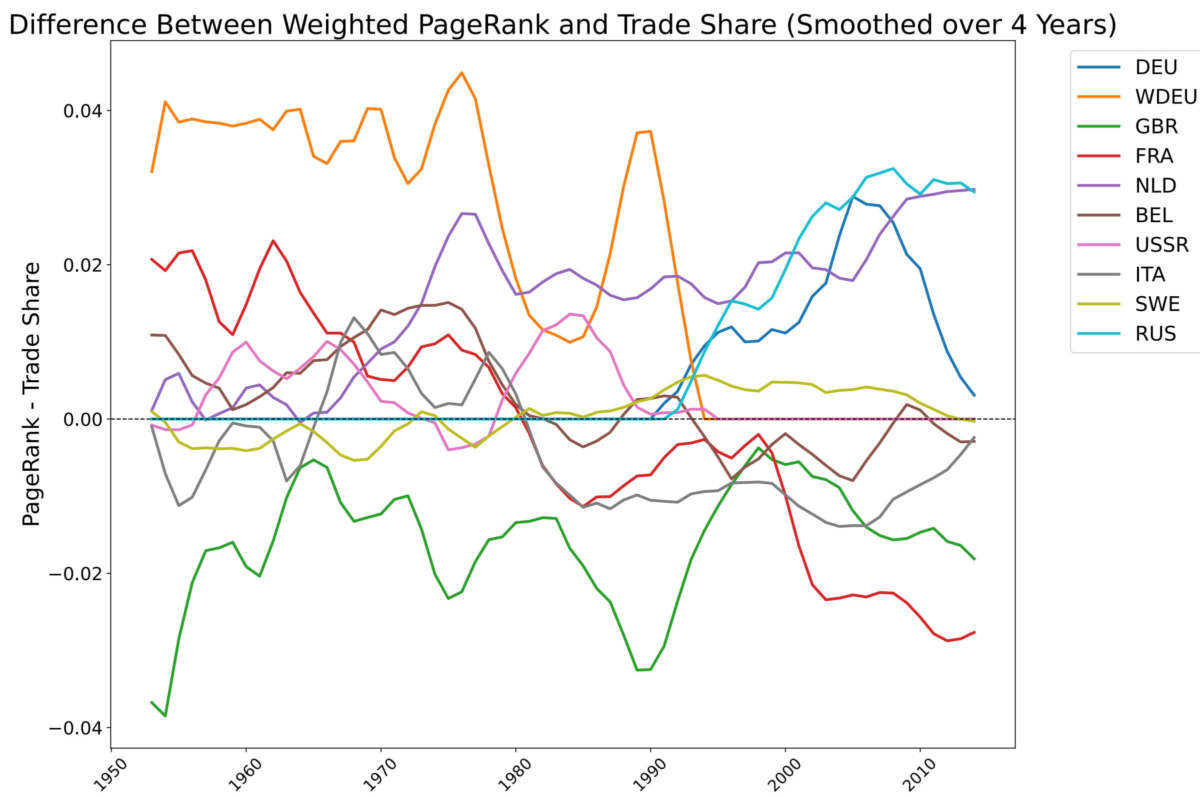

Figure 6: Weighted PageRank - Export Share for top 10 countries of the intra-European trade network. PageRanks are computed using the transpose of the weight matrix. $\alpha = 0.99$

Considering the weighted PR of countries in just the intra-European trade network, the chart of each country's weighted PageRank for each year is shown in Figure 5. Unlike in the graph of all countries, where most European states show a decline in their overall weighted PRs, when it comes to intra-European trade, most nations maintain a similar weighted PR between 1960 and 2014, apart from a decline in the UK's score. Another standout is the rise in the USSR's PR starting in the early 1980s followed by a sharp decline leading up to its dissolution. Additionally, since 1990s Russia's export PR has been increasing in the network of intra-European trade. Like with the global network of trade, we also plot the difference in each country's PR and its respective share of intra-European trade for each year in Figure 6. Again, we see some divergence between each country's PR and its share of exports on the continent. For example, the Netherlands and Russia have a higher PR than export share, perhaps resulting from the reliance on their respective agricultural and energy exports to the rest of Europe. Additionally, like with the US in the global trade network, we see an inversion for France as its PR is initially greater than its export share until the 1980s, before falling below it for the rest of the period. Additionally, the UK has a lower PR than export share for the entire period.

In addition to PageRank centrality, we also compute the principal eigenvector of the matrix $\hat{W}^\top \hat{W}$, where $\hat{W}$ is the row-normalized trade matrix with entries $\hat{W}_{ij}$ representing the share of country i's total exports sent to country j. The product $\hat{W}^\top \hat{W}$ yields a symmetric matrix. Using the principal eigenvector of $\hat{W}^\top \hat{W}$, we can assign a scalar score to each country, indicating its centrality in the trade network. This approach is inspired by the authority score from Kleinberg's HITS algorithm (1999), in which nodes that receive links from many important hubs emerge as authoritative. The application of this method for temporal networks has been discussed in Taylor et al. (2016). Like with the weighted PR score charts, we plot the authority scores of the top ten countries in the world trade network and the intra-European trade network in Figure 32 and Figure 33 respectively. For the most part, the top ten countries in terms of authority score centrality are the same as for weighted PR, however, the two measures differ in what type of centrality they capture. PageRank emphasizes recursive importance: a country receives a high score if it trades a lot with other high-ranking countries. In contrast, the authority score measures similarity rather than influence, as countries are ranked higher when their trading relationships are similar to those of other high-scoring countries, regardless of total volume of trade. As a result, countries with modest trade volume but structurally typical trade patterns may receive higher authority scores than under PageRank, and vice versa.

4.2 The Modularity and Connectivity of the Trade Network

Similar to social and biological networks, trade networks are characterized by communities or clusters, where countries tend to trade more extensively with others within the same community than with those outside of it. These trade communities may parallel the grouping of countries by identical consumption preferences, factor endowments, and proximity (Krugman, 1979). The detection of these network communities has generated an extensive field of study. In particular, methods from spectral graph theory have been used extensively in community detection problems for undirected graphs (Chung, 1997). Below, we will define the types of graph Laplacians we will be using to analyze trade networks. We will also discuss how the graph Laplacian relates to the conductance of the network, which provides a notion of how 'bottlenecked' and how connected the network is.

4.2.1 Graph Laplacians

There has been extensive literature on the study of graph Laplacians for undirected graphs and the interpretations of their respective eigenvalues (Chung, 1997; Spielman, n.d.). However, less work has been done to analyze weighted directed graphs using graph Laplacians. In this section, we will give definitions of the graph Laplacian, the normalized graph Laplacian, and a definition of the directed graph Laplacian using random walks as defined by Chung (2005). Other notions of the directed graph Laplacians exist, however, we use Chung's definition as it allows us to preserve edge weights as the probabilities of the random walks.

For a weighted undirected graph with no self-loops, $G = (V, E)$, with vertex set $V = \{1, \ldots, n\}$, the graph Laplacian is defined as,

$$L = D - W,$$

where W is the weight matrix of the edges, and $D = diag\{d_i\}$ is the weighted degree matrix where $d_i = \sum_{i \in V, i \neq j} w_{ij}$.

For the same graph, the normalized Laplacian is defined as

$$L_N \coloneqq D^{-1/2} L D^{-1/2} = I - D^{-1/2} W D^{-1/2}.$$

Since L_N is symmetric, its eigenvalues are all real and non-negative. We can find n eigenvalues (not necessarily unique), and label them $0 = \lambda_1 \leq \lambda_2 \leq ... \leq \lambda_n \leq 2$. It can be shown that the smallest eigenvalue of L_N will always be zero, λ_1 = 0, and that the upper bound on these eigenvalues is 2 (Chung, 1997). The graph G is connected if and only if $\lambda_2 > 0$. More generally, the multiplicity of the zero eigenvalue is equal to the number of connected components in G, that is, if $\lambda_i = 0$, the graph has at least i disjoint components (Chung, 1997). Lastly, in the case of a connected graph, the value of λ_2 indicates the algebraic connectivity of the graph (Bapat, 2010).

Before we define the directed Laplacian, we must first explain how to represent weighted directed graphs as a random walk. In general, a random walk on n vertices is defined by an $n \times n$ transition matrix, P, where $P(i, j)$ gives the probability of transitioning to node j from node i. When P is aperiodic and irreducible, it has a unique stationary distribution, π. There are a few ways to define $P(i.j)$ given the underlying graph. First, define S as a row stochastic matrix with entries obtained by dividing each entry of W by its respective row sum,

$$S_{ij} = \frac{W_{ij}}{\sum_{k=1}^{n} W_{ik}} \; .$$

While the matrix S, defines a valid transition matrix for a random walk on the network, we modify it to ensure that the resulting transition matrix is both irreducible (every state can be reached from any other state) and aperiodic (the greatest common divisor of return times to each node is 1). Satisfying both these conditions is important, as we will require the transition matrix to have a unique stationary distribution. To implement this—like in the computation of weighted PageRank scores—we introduce a small probability of jumping uniformly to any other node, following the standard approach. Let J be the matrix where every entry is $1/n$, i.e.,

$$J = \frac{1}{n}\mathbf{1}\mathbf{1}^\top,$$

so that $J_{ij} = 1/n$ for all i, j. The final transition matrix P is defined as a convex combination of S and J:

$$P = \alpha S + (1 - \alpha)J,$$

or element-wise:

$$P_{ij} = \alpha \cdot \frac{W_{ij}}{\sum_k W_{ik}} + \frac{1 - \alpha}{n},$$

where $\alpha \in [0, 1]$ is the damping factor. As $\alpha \to 1$, the walk closely follows the original network structure; as $\alpha \to 0$, it becomes increasingly random. In our analysis, we set $\alpha = 0.99$ to produce a transition matrix with almost the same structure as the underlying network while simultaneously ensuring the transition matrix is irreducible and aperiodic.

Since P is irreducible and has strictly positive entries, the PerronFrobenius Theorem guarantees the existence of a unique left eigenvector with strictly positive components, known as the Perron vector, denoted by ϕ. If we normalize this vector such that $\sum_{v \in V} \phi(v) = 1$, then by construction, given the irreducibility and aperiodicity of P, ϕ is equal to the unique stationary distribution of the random walk, which we denote as the probability vector, π. Thus, we have $\phi = \pi$.

Using this transition matrix P and the corresponding stationary distribution π, we now define the directed Laplacian following the formulation of Chung (2005). For a weighted directed graph G with transition matrix P and stationary distribution π, the directed Laplacian is given by:

$$\mathcal{L} := I - \frac{\Pi^{1/2} P \Pi^{-1/2} + \Pi^{-1/2} P^{\intercal} \Pi^{1/2}}{2} ,$$

where Π is a diagonal matrix with entries $\Pi(u, u) = \pi(u)$ and $P^{\intercal}$ is the transpose of P. As with the normalized Laplacian for undirected graphs, the eigenvalues of $\mathcal{L}$ satisfy

$$0 = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n.$$

Similarly, the multiplicity of the zero eigenvalue corresponds to the number of strongly connected components in the graph, and the second smallest eigenvalue, λ_2, represents the "algebraic connectivity" of the graph, capturing how well-connected the network is. Notably, if $G = (V, E)$ is an undirected graph, then this formulation of $\mathcal{L}$ reduces to the standard normalized Laplacian of P (Gleich, 2006).

We can also relate the eigenvalues of $\mathcal{L}$ to the eigenvalues of the underlying transition matrix P. Let ρ_i denote the eigenvalues of P corresponding to its left eigenvectors. Since P is irreducible and aperiodic, we have $\rho_1 = 1$, and the real parts of all other eigenvalues satisfy $|Re(\rho_i)| < 1$

(Levin et al., 2009). Using Chung's (2005) work, we can define an upper-bound for λ_2 as follows:

$$\lambda_2 \leq \min_{\rho_i \neq 1} (1 - Re(\rho_i)).$$

For brevity's sake, we will refer to $\min_{\rho_i \neq 1} \ (1 - Re(\rho_i))$ as the real part of the spectral gap of P—though this is not a formal definition. For reversible Markov chains, all eigenvalues of P are real, and λ_2 will be less than or equal to the spectral gap of P: $1 - \rho_2$. For reversible chains, the spectral gap provides a bound on the mixing time and conductance of P (Levin et al., 2009). However, less work has been done on defining and interpreting a spectral gap for non-reversible Markov chains. Chatterjee (2023) defines a spectral gap for general chains, but for this analysis, we will simply be using the inequality above, as it allows us to relate the spectral gap to λ_2 (Chung, 2005).

4.2.2 Methods to Symmetrize the Trade Network

Computing the normalized Laplacian matrix requires an undirected weight matrix for each year of trade, so computing the algebraic connectivity requires some symmetrization of the weighted directed network of trade. Fagiolo et al. (2010) discuss contributions that utilize both directed and undirected approaches to analyzing the world trade networks. They make the point that analyzing the undirected network of trade might underestimate the role of heterogeneity in bilateral trade flows. For instance, a lot of trade linkages are heavily skewed to one direction, such that intense trade linkages coexist with relatively weaker linkages. Consequently, we must consider under what conditions the network should be symmetrized, and by which method the symmetrization is computed. Fagiolo (2006) provides a procedure for deciding whether or not an observed network is sufficiently symmetric to warrant an undirected analysis. Fagiolo (2010) concludes that the international network of aggregate trade flows is sufficiently symmetric to warrant an undirected analysis. Furthermore, we would argue that in assessing the connectivity of internal trade, a symmetrization of the aggregate flows network provides some proxy for the weight of an economic relationship and interdependence between states, so the algebraic connectivity may serve to capture the overall interdependence of the network.

As for the method used to symmetrize the network, we need to consider the structural properties of the trade network we hope to preserve. If the priority is to have edge weights capture some degree of the volume of trade between edges, summing or taking the maximum of the two edge weights would be a reasonable method, i.e.

$$W_{sym}^{t} = (W + W^{\intercal}) \quad \text{or} \quad w_{sym,ij} = max\{w_{ij},\ w_{ji}\}.$$

If, instead, the priority is to capture how strong the bidirectional ties of trading relationships are, or how reciprocal they are, averaging the edge weights, or a more involved approach where we take the square root of the product of the edge weights, can be used.

For this analysis, since we are interested in the political connections associated with trade and overall network connectivity, we use an averaging of the edge weights, so our symmetric matrices are of the form,

$$w_{sym,ij} = \frac{1}{2}(w_{ij} + w_{ji}) \quad \Longrightarrow \quad W_{sym} = \frac{1}{2}(W + (W)^{T}).$$

So entry $w_{sym,ij}$ is equal to the average value of exports from country i to country j and exports from country j to country i.

To reiterate Fagiolo's point, we find that the less symmetric the initial trading networks are, the more algebraic connectivity varies based on the symmetrization method used. So, for intra-European trade, which is relatively more symmetric than other trade networks, the choice of symmetrization method has less of an impact on the algebraic connectivity of the network. In contrast, when we apply these methods to networks with much more asymmetric trading flows, like commodity trade networks, the normalized graph Laplacian varies significantly based on which symmetrization method is chosen. So, in the case of those networks, methods that allow for directed networks, like graph conductance or the directed graph Laplacian, should be used to provide a notion of network connectivity.

4.2.3 The Normalized Graph Cut

A graph $G = (V, E)$ can be partitioned into two disjoint sets, B and C, by removing edges connecting the two components. In the context of trade networks, B and C can represent two trading communities, and the degree of dissimilarity or disconnectedness between these two communities can be measured by the minimum number of edges (or sum of edge weights) that need to be removed in order to disconnect them. In the language of networks, this is called the graph *cut* and is defined as:

$$cut(B, C) = \sum_{i \in B\ j \in C} w_{ij} \ .$$

The optimal cut of a graph is thus given by the choice of subsets B and C that minimizes this cut value. Finding this minimum cut is a well-studied problem (Shi and Malik, 2000). However, one issue with just computing the minimum cut in our context is that the minimization of this value tends to favor cutting small sets of relatively isolated nodes from the graph, which we can imagine might be quite common when looking at networks of trade. To avoid this bias for partitioning out just a few countries from the network, and instead attempting to capture dissimilarity between larger trading blocs or communities, we will incorporate a measure of disassociation called the *normalized cut (Ncut)* defined by Shi and Malik (2000) as:

$$\text{Ncut}(B, C) = \frac{cut(B, C)}{assoc(B, V)} + \frac{cut(C, B)}{assoc(C, V)} \ .$$

where $assoc(B, V) = \sum_{i \in B\ j \in V} w_{ij}$ captures the total connection from nodes in B to all nodes in the network–with $assoc(C, V)$ being similarly defined. Conceptually, minimizing this problem balances finding communities that are disconnected while simultaneously trying to maintain relatively equal amounts of trade originating from each community. This ensures that a potential partition of the network does not simply cut off a small or weakly trading set of countries simply because they have little trade to begin with. Finding $B, C \subset V$ that minimizes $\text{Ncut}(B, C)$ is an NP-complete problem; however, following the approach developed by Shi and Malik (2000), rather than solving the exact combinatorial problem of minimizing the normalized cut, they instead relax the problem into a continuous optimization problem. The relaxed problem is a generalized eigenvalue problem with the solution being given by the eigenvector corresponding to λ_2 of the normalized Laplacian. This eigenvector embeds the countries into $\mathbb{R}$, where countries that are more interconnected in the graph will tend to have closer values in the vector. Using the values of this eigenvector, we can partition the countries using a variety of methods, where a common approach involves placing all countries with a positive entry in the Fiedler vector in one subset and all those with a negative entry in the other subset. The choice of subsets given by the Fiedler vector approximately minimizes the normalized cut. Beyond the entries of the eigenvector, the magnitude of λ_2 can be interpreted directly. Small values of λ_2 imply few edges need to be removed to split the network, corresponding to a more modular structure of the network with weakly interconnected trade blocs. Larger values of λ_2, by contrast, indicate more integrated trading networks where a significant number of trading relationships would need to be removed in order to partition the network into two trading communities.

4.2.4 Network Conductance

Similar to the normalized cut, graph conductance can provide of measure of how well-connected the graph is and provide a sense of how "bottlenecked" the network of trade is. For a weighted directed graph, a common approach for computing graph conductance is to define conductance based on the flow of a random walk over the network (Levin et al., 2009; Spielman, n.d.). To do this, we start by defining a Markov chain. Given a weighted directed graph $G = (V, W)$ where $|V| = n$, for nodes i and j in the graph, define the transition probability between them using the same method as for when computing the directed graph Laplacian:

$$P(i,j) = \ \alpha \frac{W_{ij}}{\sum_{k=1}^{n} W_{ik}} + \frac{(1-a)}{n} \ .$$

Since this chain is irreducible and aperiodic because of the damping factor, there must exist a stationary distribution, π, for this Markov chain. Now, for any strict subset of vertices, $S \subsetneq V$, the conductance of S, $\Phi(S)$ is defined as:

$$\Phi(S) = \frac{\sum_{i \in S} \sum_{j \notin S} \pi(i) P(i,j)}{min\{\pi(S),\ 1 - \pi(S)\}} \ ,$$

where $\pi(S) = \sum_{i \in S} \pi(i)$. This metric quantifies the probability of leaving the set S under the stationary distribution, normalized by the volume of the smaller side of the partition. The global conductance of the directed graph, Φ_G is then given by the minimum of $\Phi(S)$ over all strict subsets:

$$\Phi_G := \min_S \Phi(S),$$

and Φ_G is in the range $0 < \Phi_G < 1$. A low conductance indicates that there exists a relatively small number of edges connecting two parts of the graph together, indicating a higher modularity, where a larger conductance represents a more connected and less modular graph. Further down, for each year of trade data, we will plot the corresponding values of Φ_G alongside the values of λ_2.

Unfortunately, computing graph conductance is an NP-complete problem since testing all subsets of vertices has a time complexity of $O(2^n)$. However, we can put a bound on Φ_G using the second smallest eigenvalue of the directed graph Laplacian. Using Chung's definition of the graph Laplacian for a directed graph, $\mathcal{L}$ with eigenvalues ($\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$), Chung proves that Cheeger's inequality,

$$2\Phi_G \geq \lambda_2 \geq \frac{\Phi_G^2}{2} ,$$

still holds for the eigenvalues of $\mathcal{L}$ just as it does for the eigenvalues of L_N (Chung, 2005). Since we want to bound Φ_G, we can rearrange the above inequality to get

$$\frac{\lambda_2}{2} \leq \Phi_G \leq \sqrt{2\lambda_2} .$$

This inequality is tight, however, once we fix G, we can obtain a potentially better upper bound on Φ_G by computing $\Phi(S)$ for sets that we expect to have a small conductance, and then the minimum of those computed values can be used as a new upper bound. To do this, we use an approach that involves sweeping through the Fiedler vector to select subsets.

This method begins with computing the directed Laplacian, $\mathcal{L}$ of the network, and the eigenvector corresponding to λ_2, the Fiedler vector, $\mathbf{v_2}$. Each country, i has a corresponding entry in the Fiedler vector, $\mathbf{v_2}[i]$. By sorting countries based on their Fiedler vector values, we obtain an ordered list of countries $(v_1, v_2, \ldots, v_n)$. From this ordering, we define a sequence of nested subsets:

$$S_k = \{v_1, v_2, \ldots, v_k\}, \qquad k = 1, 2, \ldots, n-1,$$

and

$$S'_k = \{v_n, v_{n-1}, \ldots, v_{n-k+1}\}, \qquad k = 1, 2, \ldots, n-1,$$

For each S_k we compute $\Phi(S_k)$, and for each S'_k we compute $\Phi(S'_k)$, and then take the minimum value across all k as our approximation of Φ_G. Since Φ_G is defined as the minimum conductance over all subsets, the process gives a valid upper bound on the true graph conductance. In the next section, we will plot λ_2, the bounds given by Cheeger's inequality, and our approximated Φ_G for each network over time.

4.3 Comparative Theorems for Graph Laplacian

Having discussed how the second smallest eigenvalue of the Laplacian, λ_2, reflects the connectivity and modularity of the trade network, we move now to discuss how λ_2 is expected to change over time as the network of trade evolves. First, we discuss how changes in the overall magnitude of trade flows across the network affect λ_2. Second, we explore how the introduction or removal of countries in the trading network alters λ_2.

4.3.1 Uniform Scaling of the Trade Network

We can consider an increase in trade across the network as a perturbation of the weighted adjacency matrix representing exports between countries. Consider a network of trade, $G = (V, E)$, where $W \in \mathbb{R}^{n \times n}$ is the weighted adjacency matrix of the network. As above, we can construct a transition matrix, P using this network, and a corresponding directed Laplacian, $\mathcal{L}$, with second smallest eigenvalue, λ_2. Now consider a new network given by a scaling of W by a constant factor, $\gamma > 0$, as $\tilde{W} = \gamma W$. If we were to construct the transition matrix of this new network in the same method as before, we would have.

$$\tilde{P}_{ij} = \alpha \frac{\gamma\, w_{ij}}{\sum_{k=1}^{n} \gamma\, w_{ik}} + \frac{1-\alpha}{|V|} = \alpha \frac{w_{ij}}{\sum_{k=1}^{n} w_{ik}} + \frac{1-\alpha}{|V|} = P_{ij}\ .$$

And since the transition matrices are the same, $\tilde{\mathcal{L}} = \mathcal{L} \implies \tilde{\boldsymbol{\lambda}}_2 = \boldsymbol{\lambda}_2$. That is, the directed Laplacian does not change under a linear scaling of the weighted adjacency matrix, and thus the connectivity of the network does not change under a uniform linear scaling of trade volumes.

4.3.2 Disruption of Individual Trading Relationships

We now consider how the removal of a single trade relationship affects the overall connectivity of the trade network. Specifically, we examine how such disruptions influence the graph's algebraic connectivity, as measured by the second smallest eigenvalue λ_2 of the normalized Laplacian. While the true impact of a disrupted trade relationship depends on both the magnitude of the trade and the broader structure of the network, we can derive upper and lower bounds on the new value of λ_2 after the removal. Of course, a major limitation of these theoretical bounds is that they assume that the rest of the network structure remains fixed after the disruption, as we do not model any re-optimization or trade diversion. So the benefit of this analysis is not to produce an impulse response function, but instead to determine the direction of bias for the change in λ_2 following a shock.

Consider a network of trade, $G = (V, E)$ where $W \in \mathbb{R}^{n\times n}$ has been symmetrized and has normalized Laplacian, $L_N(G)$. Then let $G - e$ be a subgraph of G where one edge, $e \in E$, is removed. This could represent, for instance, the severing of a bilateral trade link due to sanctions or political breakdown. Let $\lambda_1 \leq \cdots \leq \lambda_n$ and $\mu_1 \leq \cdots \leq \mu_n$ denote the eigenvalues of the normalized Laplacian of G and $(G - e)$ respectively. Then the following interlacing inequality holds (Chen et al., 2004; Butler, 2007):

$$\lambda_{i-1} \leq \mu_i \leq \lambda_{i+1}, \qquad \forall i \in \{1, \ldots, n\}.$$

There may exist a further restriction of the above inequality that allows us to compare the algebraic connectivity of G and $(G - e)$—as such an inequality exists for non-normalized Laplacians (Chen et al., 2004). Finding a similar inequality for the normalized Laplacian would assist with analyzing simulated trade networks. While not proven, intuitively, as individual trade relationships are disrupted, we would expect an overall more modular network.

4.3.3 Removal or Addition of a Country in the Network

To interpret the impact of a decrease in the number of countries on the modularity of the trade network, we can consider a contraction of two countries into one trading partner. To represent this contraction as a network, we will use a formalization given by Chen et al. (2004). Consider a network of trade, $G = (V, E)$, where the weight matrix, W, has been symmetrized. Then, for any two vertices u and v of G, denote $G/\{u, v\}$ as the graph obtained by contracting (combining) u and v into a single node. The neighborhood of the new node corresponds to the union of the original neighbors of u and v. Formally, the export edge weights from the new node to a third country $i \neq u, v$ are given by $w_{ui} + w_{vi}$, and the import edge weights from i to the new node are $w_{iu} + w_{iv}$. By treating u and v as a single equivalence class while all other nodes remain as singletons, we can apply interlacing results from Butler (2007) and Aliniaeifard et al. (2021) to compare the spectra of the normalized Laplacians of G and $G/\{u, v\}$. If the eigenvalues of normalized Laplacian of G and $G/\{u, v\}$ are $\lambda_1 \leq \cdots \leq \lambda_n$ and $\nu_1 \leq \cdots \leq \nu_{n-1}$, then we have the following interlacing relationship:

$$\lambda_1 \leq \nu_1 \leq \lambda_2 \leq \ldots \lambda_{n-1} \leq \nu_{n-1} \leq \lambda_n.$$

This implies that in the event of a removal of a country from a network of trade, assuming some reallocation of the trading being imported and exported by said country, we would expect a decrease in the modularity of the overall network—assuming such a removal results in no other changes across the network. Conversely, we can reinterpret this result to consider the change in connectivity as a result of adding a country to the network. Suppose $G/\{u, v\}$ represents the current trade network, and a new country is introduced by splitting an existing node into two countries u and v where the existing export and import values are distributed between u and v, then we would have the graph, G as described above. Thus, using the same interlacing inequalities, absent any other changes across the network, if we were to add an additional country to the network, we would expect a decrease in network connectivity.

While the assumption that the rest of the network remains unaffected by the addition or removal of countries is admittedly unrealistic, theoretical comparative statics still offer useful insights into the direction of bias in network modularity. For instance, if new countries are introduced into the network without significantly altering overall trade volumes, we would expect an increase in network modularity. However, if modularity instead decreases, this suggests that other structural changes in the network are at play. In that case, the addition of the new country alone would have biased the network towards greater modularity, meaning the observed decrease must be a result of other shifts in trade relationships. For example, if countries are added to the network without a major change in trade volumes across the network, we would expect an increase in network modularity. If, instead, we observed a decreased modularity of the new network, we can conclude that this change in modularity must in some part be a function of some other shift in the network, as just the addition of the node alone would be biasing the network towards being more modular.

4.4 Observed Connectivity of World and European Trade Networks

For each year from 1950 to 2014, the five smallest eigenvalues of the directed Laplacian, $\mathcal{L}^t$, created using each year's observed global trade network, are plotted in Figure 7. In this graph, years with lower values of λ_2^t indicate that in those years, the network of trade has a lower connectivity, and that the countries could be partitioned into two disjoint trading communities through the removal of relatively few edges. In contrast, years with larger values of λ_2^t indicate years where trade is more interconnected, and more trading relationships would need to be disrupted to bi-partition the network of trade. Interpreting the results, we fail to observe any major shifts in λ_2^t since 1961, other than a gradual decline. What stands out is the sharp decline and then rebound in λ_2^t between 1951 and 1961. Significantly more work would need to be done to properly identify the causation of such a sharp decrease in trade connectivity. However, in looking at the data, this period coincides with a sharp increase in Russian exports to a few Eastern European countries. This increase in exports might explain why we see the same sharp decline and rise when we look at just

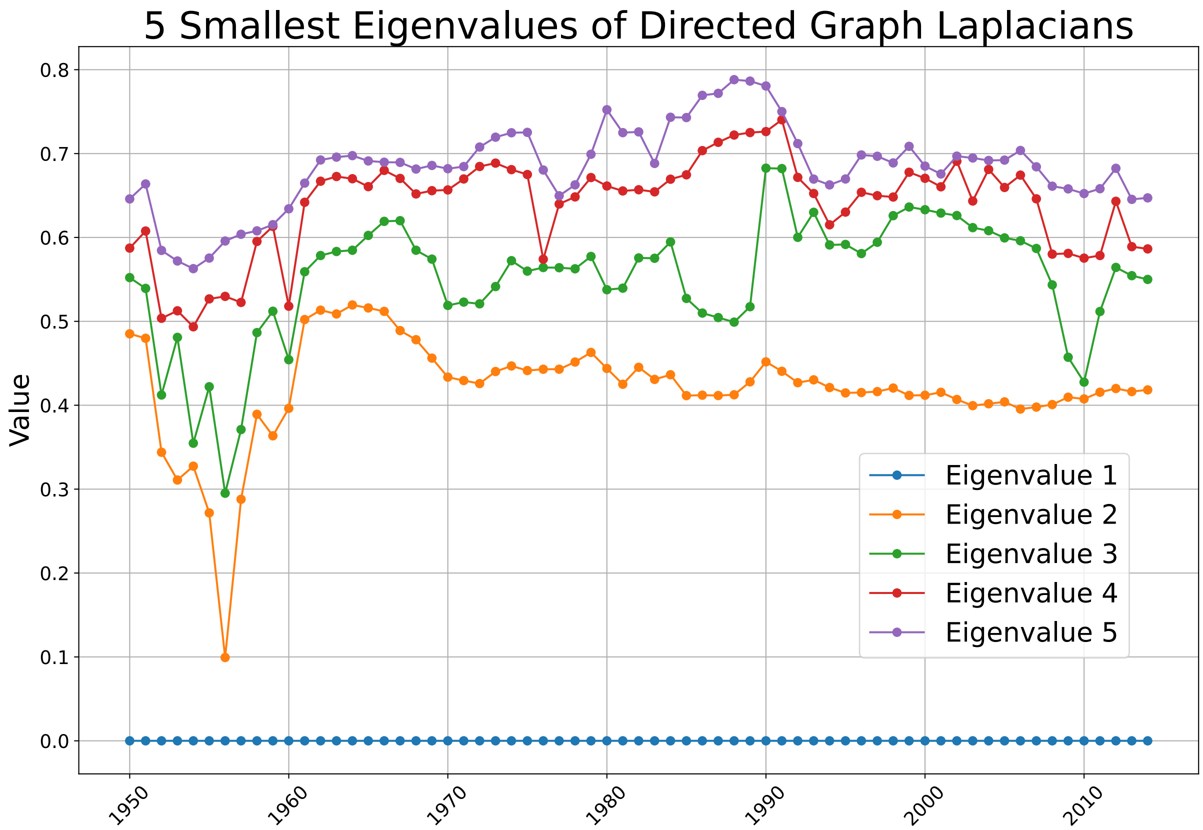

Figure 7: The five smallest eigenvalues of $\mathcal{L}^t$. λ_2^t in orange.

intra-European trade further down. While we are mainly concerned with the value of λ_2^t in this analysis, it is also important to consider the full spectrum of Laplacian eigenvalues as these can be used to interpret multi-community separations and local bottlenecks in the graph. For example, in the above chart, we observe a spike up in λ_3^t between 1989 and 1990, despite an otherwise minor increase in λ_2^t. Again, while we cannot make any claims about causality, we do observe a similar spike in λ_2 when we look at just the network of intra-European trade. So perhaps this change in λ_3^t is capturing a greater connectivity within a trading community identified by λ_2^t.

In looking at the intra-European trade, the five smallest eigenvalues of $\mathcal{L}^t$ for each year are shown in Figure 8. Similar to the world trade network, in the intra-European trade network, we again observe a sharp decline and rebound of λ_2^t between 1951 and 1961. When looking at the data for intra-European trade, the largest shift in flows comes from the USSR during this period, so perhaps this shift is correlated with the USSR trading substantially more with Eastern European countries

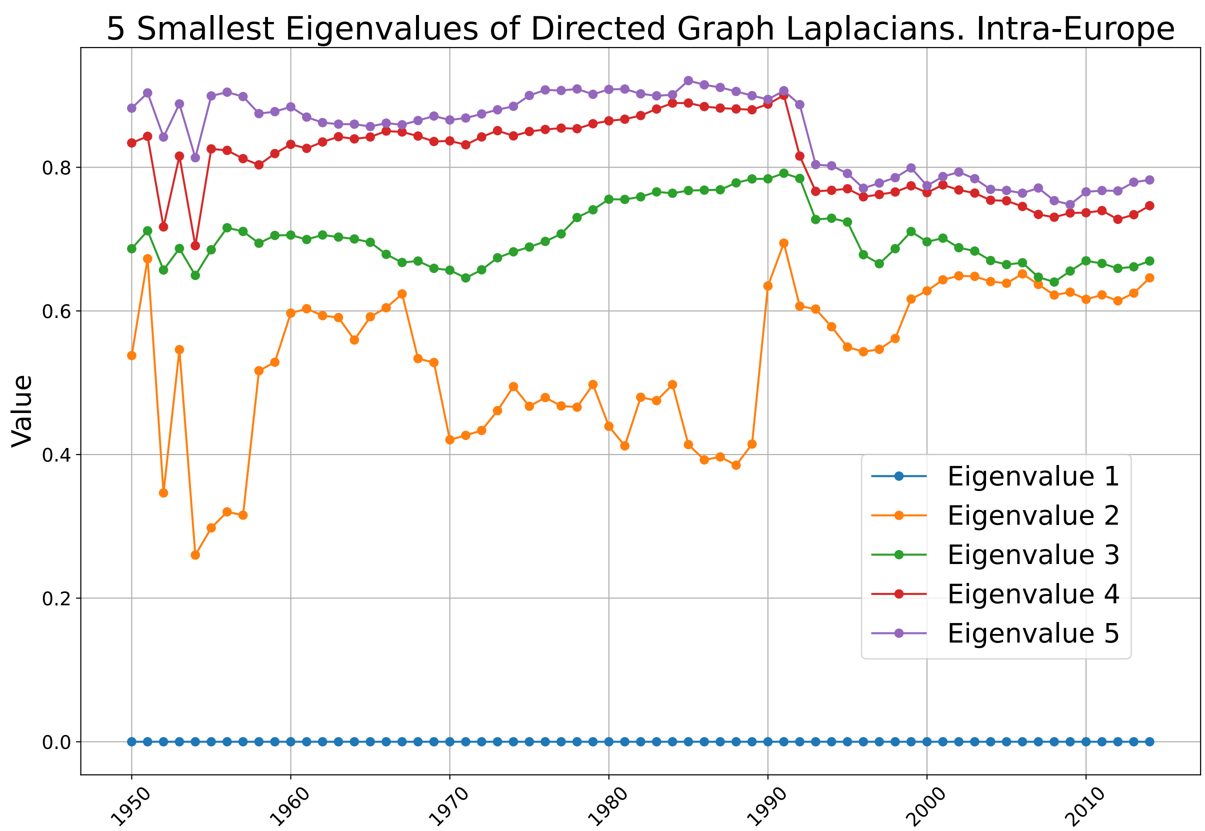

Figure 8: The five smallest eigenvalues of $\mathcal{L}^t$ for intra-European trade. λ_2^t in orange.

and then decreasing relative to overall trade on the continent. Unlike in the world trade network, λ_2^t fluctuates between 1960 and 1989 before spiking up 1990. While we make no causal claims, this increase in λ_2^t coincides with major political and economic shifts on the continent brought on by the fall of the Berlin Wall in 1989, the reunification of Germany in 1990, and then the dissolution of the USSR in December of 1991. The latter of which led to the addition of numerous new countries to the network beginning in 1992. As we discussed in the previous section, if we were to add additional countries to the network using existing trade flows, we would expect a decrease in λ_2^t, ceteris paribus. Yet the fact that λ_2^t is higher in the 1990s than it is in the 1980s indicates that shifts in the network structure other than the addition of new countries resulted in a trade network with greater connectivity.

Beyond just looking at changes in λ_2^t for the world and intra-European trade networks, we also consider how the global graph conductance, Φ_G, and the real part of the spectral gap of P change over the same time period. For both networks, we plot the following quantities: (1) λ_2 of the directed Laplacian; (2) an approximation of Φ_G using the Fiedler vector sweep method; (3) $\min_{\rho_i \neq 1}(1 - Re(\rho_i))$ which we interpret as the real part of the spectral gap of P; and (4) the theoretical bounds on the actual value of Φ_G provided by Cheeger's inequality:

$$\frac{\lambda_2}{2} \leq \Phi_G \leq \sqrt{2\lambda_2}\ .$$

We plot these values for the world trade network in Figure 9 and for the intra-European network in Figure 10. As discussed in the previous section, we can use the approximated value of Φ_G as a new upper bound on the actual value Φ_G. Charts using the approximated value of Φ_G as the new upper bound on Φ_G for the world and intra-European networks are shown in Figure 34 and Figure 35, respectively, in the appendix.

In interpreting these charts, it is important to keep in mind that we cannot directly compare specific values of λ_2 to approximated values of Φ_G. However, we can interpret their co-movements or lack thereof. For the world trade network, the approximated value of Φ_G lines up quite closely with movements in λ_2 for the whole period. For the intra-European network, λ_2 and approximated Φ_G follow similar movements apart from a few years where they diverge. For example, from 1989 to 1990, the approximated value of Φ_G increases less sharply than the values of λ_2, which might indicate that we are over-estimating the increase in network connectivity in that year. Also, it is interesting to note that Φ_G decreases from 1998 to 2000, while λ_2 increases during that period. While we can speculate as to what causes these divergences, further work would need to be done in order to interpret what differences in Φ_G and λ_2 mean for the structure of the network. Additionally, for all years in both networks, λ_2 and $(1 - Re(\rho_2))$ are nearly equal. Mathematically, we have little to add on why that might be the case. For work on relating the eigenvalues of normalized and directed Laplacians to the eigenvalues of Markov chains, see Seabrook and Wiskott (2023). For the purposes of analyzing trade networks, computing $\min_{\rho_i \neq 1}(1 - Re(\rho_i))$ may be easier than computing $\mathcal{L}$ and λ_2 for larger networks, and thus might serve as a quicker way to compute a metric of network connectivity.

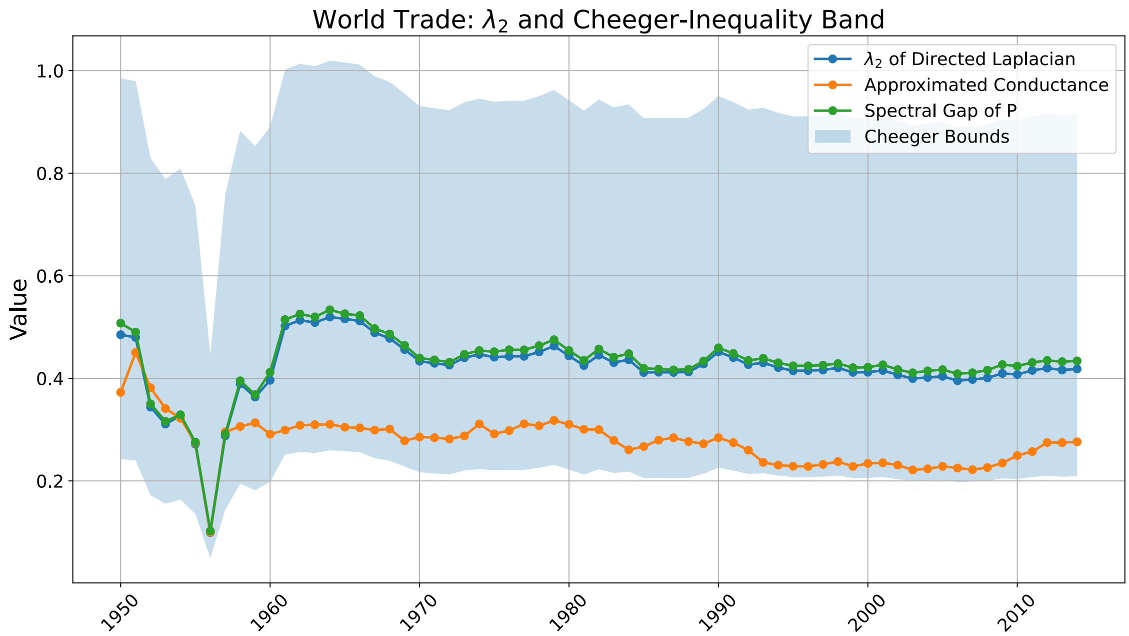

Figure 9: λ_2 of directed Laplacian, approximated Φ_G, "real part of P's spectral gap": $\min_{\rho_i \neq 1}(1 - Re(\rho_i))$, and bands on Φ_G given by Cheeger's inequality.

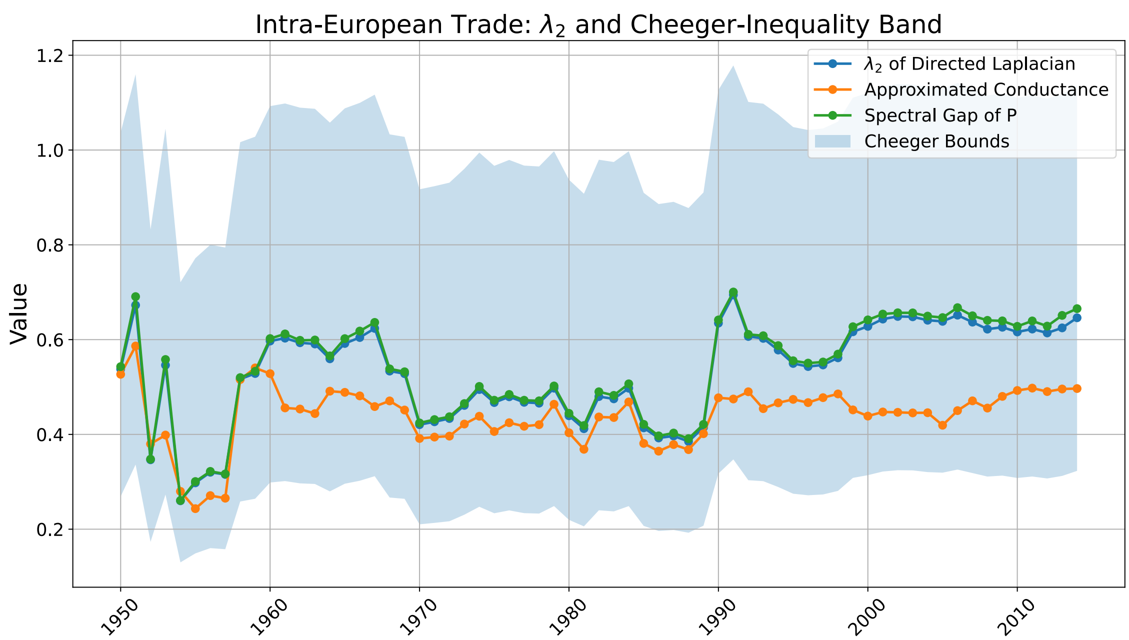

Figure 10: λ_2 of directed Laplacian, approximated Φ_G, "real part of P's spectral gap": $\min_{\rho_i \neq 1}(1 - Re(\rho_i))$, and bands on Φ_G given by Cheeger's inequality.

4.5 Comparison with Null Models

To assess the significance of the eigenvalues of the directed Laplacian, we must ask whether these values reflect meaningful structural features of the international trade network—or whether similar eigenvalues would arise in any randomly generated graph with comparable characteristics. In other words, do the observed levels of modularity and connectivity respond to real-world economic and political shifts, or would we expect to see similar eigenvalues for any random network with certain characteristics similar to the trade network? To answer this, we turn to the framework of null models—which in our analysis will be randomly generated graphs that preserve selected features of the observed empirical network (Squartini et al., 2014; Fagiolo, 2016). Specifically, we ask: given each country's total export volume and its propensity to trade more with certain partners, does the observed modularity of the trade network differ significantly from that of a random network with similar constraints? To compare how our observed graphs differ from the null models, we compare the observed eigenvalues of the trade networks to the distribution of eigenvalues that result from running a Monte Carlo simulation with 1000 replicates for a random graph that has similar characteristics to the actual trade network we observe for each year.

To construct the random graphs used in our null models, we begin by fixing every country's weighted out-degree, thereby maintaining the distribution of total export amounts in the network. For each country, i, we also fix the distribution of weights, W_i, assigned to its outgoing edges. We then randomly reassign these weights to different destination countries in the network through a surjective mapping that avoids creating any self-loops or multiedges. This can be thought of as randomly shuffling the destination of a country's exports while maintaining the same number of exporting partners and the same distribution of values exported across partners. Let the randomized graph for year t be denoted as, $\mathcal{G}_t(N_t, \{d_i\}, \{W_i\}, \mathcal{R})$, where N_t is the set of all countries in the observed network of trade for year t, $\{d_i\}$ is the set of out-degrees for each node, where d_i represents the number of outgoing edges for node i in the observed network of trade, and $\{W_i\}$ defines a set of weight vectors for each node i of the $|N_t|$ nodes in the graph. That is, for each i,

$$W_i = \{w_{i1}, w_{i2} ... w_{id_i}\},$$

where each w_{ij} denotes the weight of the outgoing edge from node i to node j in the actual trade network (which excludes self-loops from country i to itself). Lastly, we define $\mathcal{R}$ as a random permutation, or matching function, that ensures a one-to-one random assignment of the d_i outgoing edges of each node i to d_i distinct nodes in N_t. The randomness of $\mathcal{R}$ can be described as $\mathcal{R} \sim \mathcal{U}(\mathcal{P})$, where $\mathcal{P}$ is the set of all possible permutations of edge weights to edges, excluding the edges between countries and themselves.

For both the world trade network and the network of intra-European trade, we simulate 1000

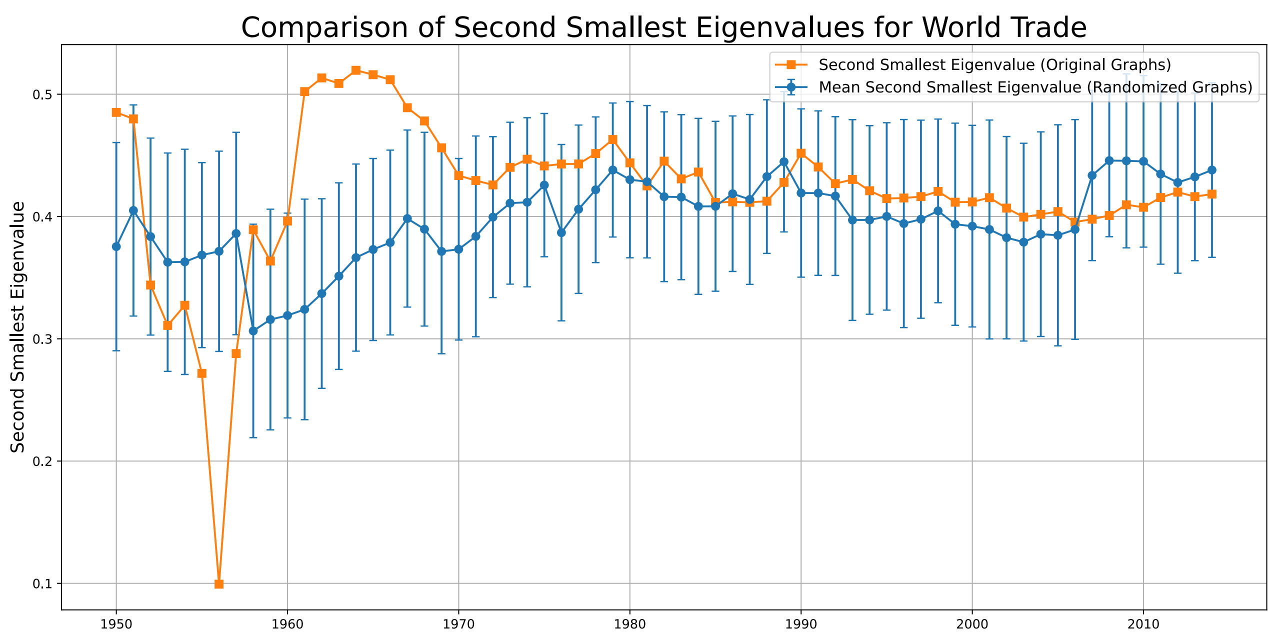

Figure 11: Second smallest eigenvalue of the directed Laplacian (λ_2) of the world trade network in each year (orange) plotted against the distribution of the second smallest eigenvalue across 1000 replications of the random graph using parameter values from the corresponding world trade network. Error bars represent two standard deviations.

random graphs for each year between 1950 and 2014 using the method described above. For each simulated random graph, we compute the directed Laplacian and the corresponding 5 smallest eigenvalues. The mean value of each eigenvalue, and the associated standard deviation across 1000 replications for each year, is plotted for the world trade network in Figure 36, and the intra-European network in Figure 37. Additionally, we plot the distribution of the second smallest eigenvalues across the replications of the random graph against the actual values of the world trade network's second smallest eigenvalue in Figure 11, and for the intra-European network in Figure 12.

For the world trade network, the second smallest eigenvalue, λ_2^t, is close to the mean value of

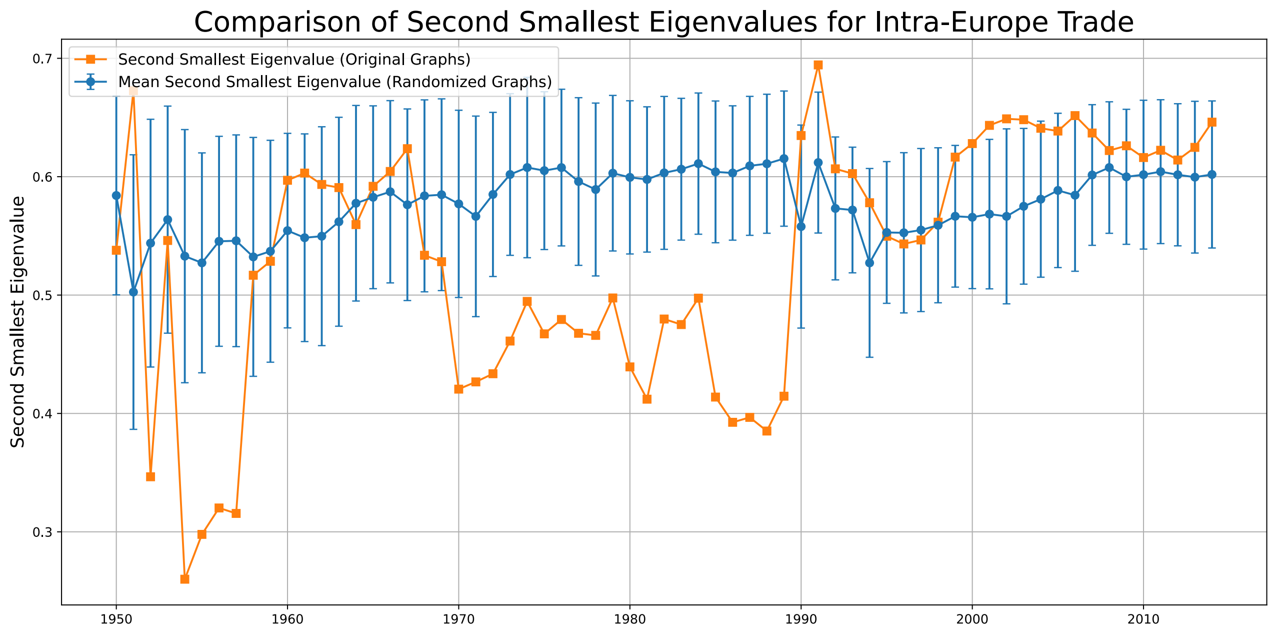

Figure 12: Second smallest eigenvalue of the directed Laplacian (λ_2) of the intra-European trade network in each year (orange) plotted against the distribution of the second smallest eigenvalue across 1000 replications of the random graph using parameter values from the corresponding intra-European trade network. Error bars represent two standard deviations.

λ_2 across 1000 simulations of the null model for most years, apart from in 1957, where the observed value of λ_2 is far below the null model distribution. The latter would indicate that the trade network was far more modular than we would expect if the network were randomly constructed. For 1961-1968, the value of λ_2 for the observed network is more than two standard deviations above what we would expect for a null model network. However, since λ_2 is for the most part close to the mean of the null model distribution, which means we may be unable to identify any special structure in the network of trade with respect to changes in connectivity.

The intra-European trade network, by contrast, exhibits stronger and more consistent deviations in network connectivity. From 1970 to 1989, the observed λ_2^t is more than two standard deviations below the mean of the null distribution. This indicates that the trade network was significantly more modular than would be expected under a random configuration during this period, suggesting that political and economic forces shaped a deliberately segmented and relatively disconnected network structure. This, of course, lines up with the economic and political impact of the Iron Curtain on Europe. Notably, in 1990, λ_2^t for the observed network spikes upward, even as the mean λ_2 from the null distributions declines. This indicates that not only is the observed network more connected than a random configuration, but the network became more interconnected in a year we would have otherwise expected λ_2 to decrease, given the distribution of countries and edge weights.

These findings are important not because they reveal anything we did not already know about the evolution of trade in Europe, but because they validate the use of λ_2 as a meaningful metric of trade connectivity. Since the observed values of λ_2 deviate significantly from the null distribution during major political and economic transitions demonstrates that λ_2 responds to real-world structural shifts, rather than just reflecting noise in the network data.

4.6 Comparison with Inter-state Trade Within the US

In addition to a comparison with null models, we can also consider the eigenvalues of the directed Laplacian we observe in already integrated trading communities, such as trade within the United States. Using data from the United States Department of Transportation. Bureau of Transportation Statistics Commodity Flow Survey (CFS) (1993, 1997, 2002, 2007, 2012, and 2017), we construct a network of state-to-state trade within the United States for six years across a 15-year range (1993-2017). While the CFS data is far from complete in terms of capturing all trade between US states, as long as the trade values reported by CFS are proportional to the actual values of trade, then as we discussed in section 4.3.1, the resulting eigenvalues of the directed Laplacian will be the same for the CFS data as they are the actual data. While we doubt that the networks constructed using the CFS data will be exactly proportional to the actual network, the scale invariance of the directed Laplacian will hopefully mitigate the impact of under-reported trade values on modularity. The five smallest eigenvalues of the directed Laplacian of inter-state US trade for each year are shown in Figure 13. In this chart, we observe that the algebraic connectivity of intra-US trade varies in a similar range as the algebraic connectivity of intra-European trade starting in 1992. While we do not claim that intra-European trade and intra-US trade are directly comparable in all respects, this analysis provides valuable context for understanding how trade modularity evolves in integrated economies. The fact that algebraic connectivity levels converge to a similar range following the year 1992 suggests that, at least in terms of network structure, Europe's trade integration process brought it closer to resembling an internal US trade network rather than a set of distinct, self-contained regional economies.

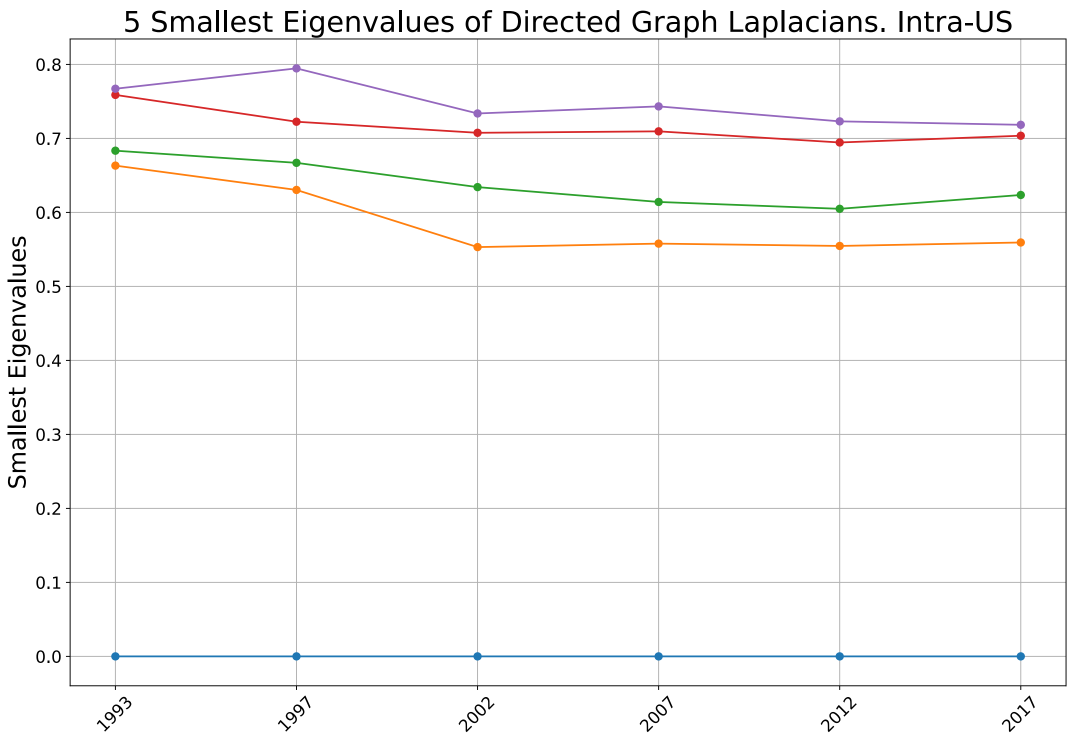

Figure 13: Five smallest eigenvalues of the directed Laplacian of inter-state US trade. λ_2^t in orange.

4.7 Partitioning Trade Networks and Community Detection

Beyond expressing a degree of modularity in the network, we can also use the directed Laplacian to explore the formation and evolution of trade blocs in the trade network. In particular, we proceed now to explore the evolution of trade blocs in the intra-European network of trade before and after the dissolution of the USSR. Earlier, we introduced the normalized cut (Ncut) as a method for partitioning the graph in a way that minimizes the total weight of edges crossing between two subsets while maintaining relatively balanced trade volumes on each side.

Given the directed Laplacian matrix $\mathcal{L}$ for each year, we compute the eigenvector corresponding to the second smallest eigenvalue λ_2^t, known as the Fiedler vector (Fiedler, 1973; Wyss-Gallifent, 2019). We denote this vector as v^t. Each entry in v^t corresponds to a country in the network. Following the relaxation of the Ncut problem, we assign countries to trade communities based on the value of their corresponding entry in v^t (Shi and Malik, 2000). There are a variety of clustering algorithms that we can use to divide the entries of the Fiedler vector, such as k-means or hierarchical clustering (von Luxburg, 2007). In our analysis, we will use hierarchical clustering as this method allows us to visualize sub-communities on both sides of the partition, and it visualizes a distance of dissimilarity between the partitions. In the following charts, we use this spectral clustering method to visualize the reorganization of intra-European trade communities before and after the dissolution of the USSR.

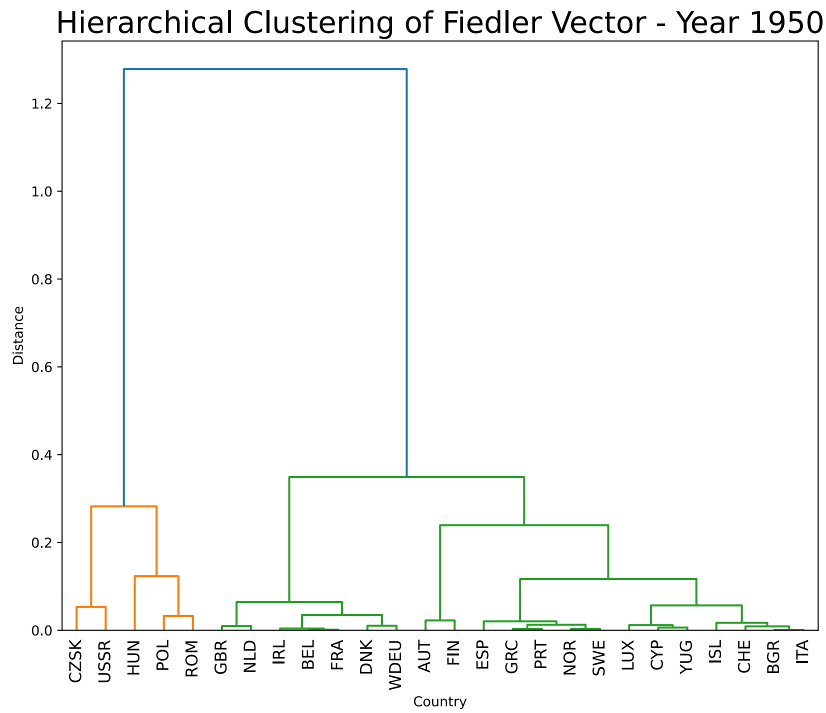

Figure 14: Hierarchical Clustering using λ_2 of $\mathcal{L}$ of the intra-European trade graph. 1950.

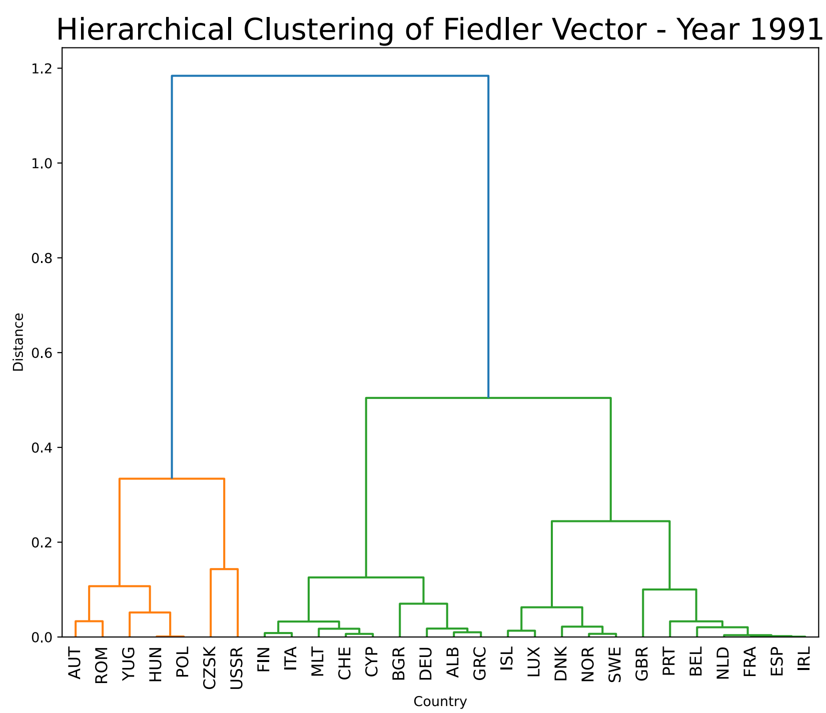

Figure 15: Hierarchical Clustering using λ_2 of $\mathcal{L}$ of the intra-European trade graph. 1991.

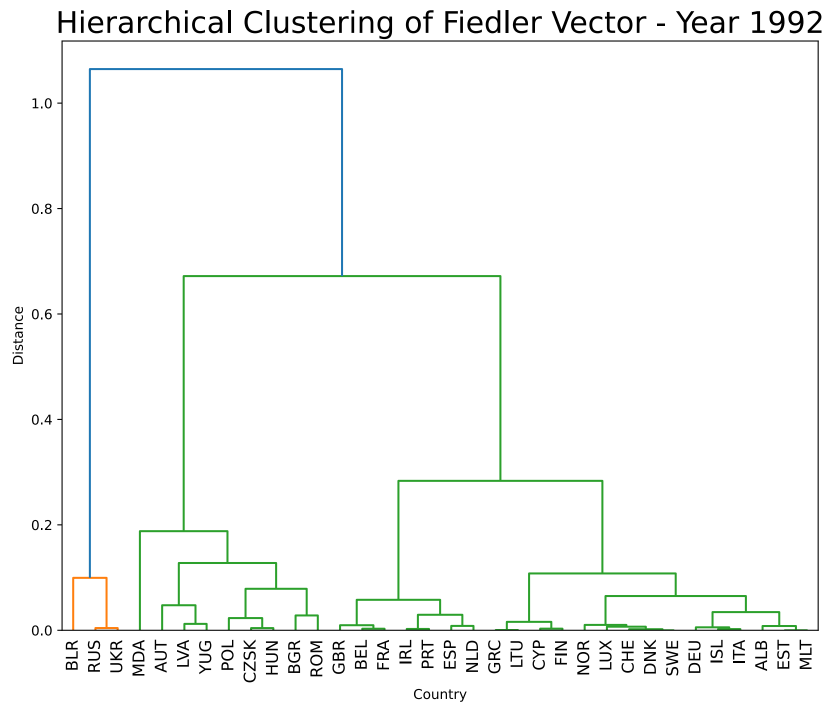

Figure 16: Hierarchical Clustering using λ_2 of $\mathcal{L}$ of the intra-European trade graph. 1992.

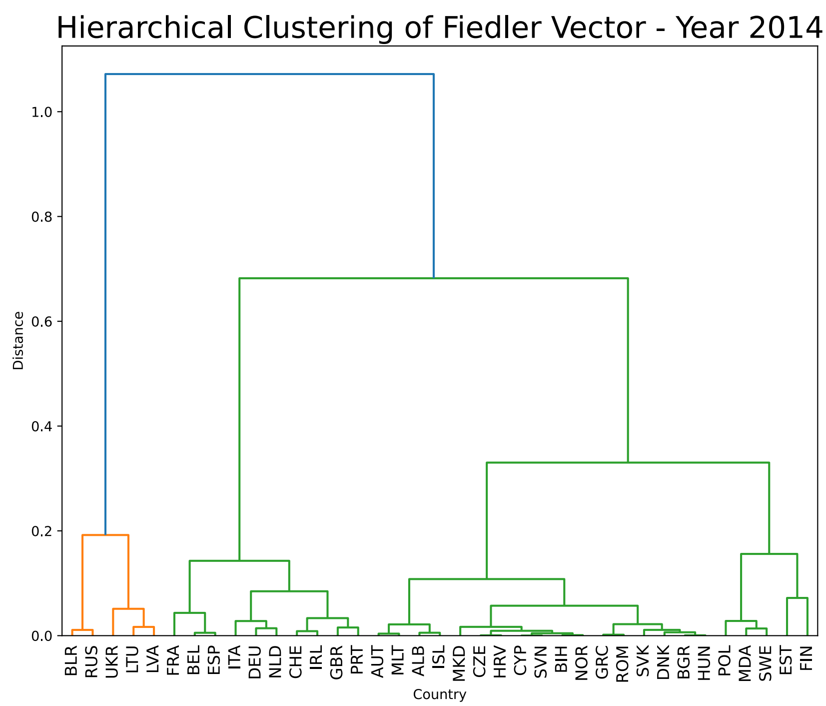

Figure 17: Hierarchical Clustering using λ_2 of $\mathcal{L}$ of the intra-European trade graph. 2014.

Hierarchical clustering constructs a tree-like hierarchy of data points by iteratively merging or splitting clusters based on a chosen distance metric. We run hierarchical clustering on v^t using Ward's method to compute the distance between clusters (Ward, 1963). The results of hierarchical clustering for the intra-European trade graph are shown for the years 1950, 1991, 1992, and 2014 in Figures 14, 15, 16, and 17, respectively. In interpreting these partitions, the main metric is the distance between the final linkage (where the two partitions connect at the top of the graph) and the last sub-linkage within one of the two partitions. The greater this distance, the greater the separation between the partitions. Maps based on the partitions are displayed for 1991 in Figure 38 and 1992 in Figure 39.

Looking at the hierarchical clustering for 1950, we see the Eastern Bloc countries of Czechoslovakia, the USSR, Romania, Hungary, and Poland split off into their own sub-community separate from the rest of European trade. This same Eastern Bloc cluster still exists in 1991, and is joined by Yugoslavia and Austria, both of which were on the other side of the partition in 1950. This may be a result of Austria's successful trade relationship with the USSR and satellite states, following Austria's independence from occupation in 1955 (Sorokin, 2022). While the partition between 1950 and 1991 remained similar, the distance between the last sub-linkage and the final linkage was shorter than in 1950, indicating that the two trade communities were more connected in 1991 than in 1992.

With the dissolution of the USSR in December 1991, the resulting partition of intra-European trade is quite different in 1992. The breakup of the USSR added Belarus, Moldova, Russia, Ukraine, and the Baltic states to the network. The resulting partition isolates Belarus, Russia, and Ukraine into a community separate from the rest of Europe. However, the distance between the final linkage and the last sub-linkage is much smaller than in 1991, indicating that overall, following the collapse of the Eastern Bloc, the trade network became less modular. It is interesting to note that in 1992, for all the countries previously in the Eastern Bloc, apart from the aforementioned three, these countries formed a new sub-community on the other side of the partition, as indicated by the distance between the penultimate linkage and the one below it on the graph. This is further supported by the fact that there was no significant change in the third smallest eigenvalue of $\mathcal{L}$, λ_3^t, between 1991 and 1992, which suggests that the ability to divide the graph into three clusters remains similar despite the dissolution of the USSR. In 2014, Belarus, Russia, and Ukraine were still in their own community, just now joined by Latvia and Lithuania. However, the distance before the final linkage is again smaller than it was in 1992, indicating a further decrease in the network's modularity. The sub-community of the other Eastern Bloc countries that was visible in 1992 is no longer present in 2014. In further work, the temporal evolution of sub-communities like these could be analyzed using a broader spectrum of the directed Laplacian's eigenvalues beyond just λ_2.

5 Trade in Products and Product Categories

In this section, we briefly explore the application of the above methodology to more granular trade networks, such as networks of trade in products and product categories. To this end, we will be using the aforementioned Gaulier and Zignago (2025) data, in which they have compiled bilateral trade flows for more than 200 countries and 5000 products as categorized by Harmonized System Codes (HS codes) for the years 1995-2023. Using this data, for each product (or product category), h, we construct panel network data where each year is represented as a weighted directed graph, where edge weights $w_{ij}(h)$ represent the total nominal value of product h exported from country i to country j in year t. Below, we aggregate flows to the first two digits of HS codes for product categories, or the first four digits of HS codes for products, and proceed to analyze a few example product networks using the same methodology as before.

5.1 Trade in Integrated Circuits

Aggregating flows to the first four digits of HS codes, we keep all flows labeled with HS code 8542, which represents trade in integrated circuits and closely related input components. Using this subset of data, we compute the weighted exporting PageRank (PR) of all countries exporting integrated circuits for each year of data. The results are presented in Figure 18, where we plot the yearly evolution of PR scores and label the ten countries with the highest average PR over the sample period. For each of these top ten countries, we plot the difference between each of these countries PR score and their share of total exports for HS 8542 in each year (i.e., their weighted out-degree) in Figure 40. This comparison allows us to assess whether countries are more central in the trade network than their raw export volumes would suggest, or conversely, whether they are less central despite being major exporters. In interpreting the results of these two figures, two things are clear. First, we observe a decline in the centrality of the United States and Japan as they go from the most central countries in the network to below the top five most central countries in 2023. During the same period, we observe the rise of Taiwan, China, and Korea as central countries as their respective PR scores increase. Second, in looking at the difference between PR and exports, we see that Taiwan appears substantially more central in the network than its export volume alone would imply, while China, despite being one of the largest exporters, has a relatively lower PR score. This difference is potentially capturing the global significance of Taiwan in this industry, versus a more regional significance of China.

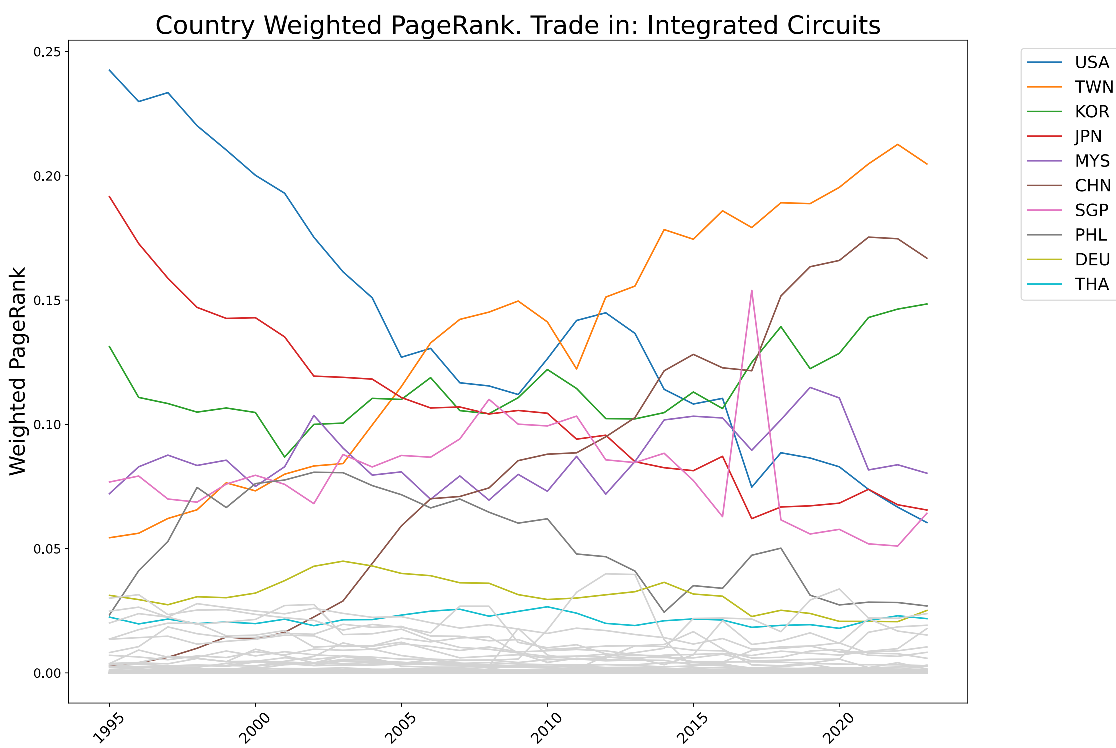

Figure 18: Weighted PageRank scores shown for countries in trade of integrated circuits (8542). Transpose of the weight matrix used to identify exporters. $\alpha = 0.99$

In looking at the evolution of connectivity in the trade of integrated circuits, we compute the directed Laplacian, $\mathcal{L}$, for each year of trade, and plot its five smallest eigenvalues in Figure 19. In this chart of eigenvalues, we observe a slight downward trend in the value of λ_2 from 1995 to 2023,

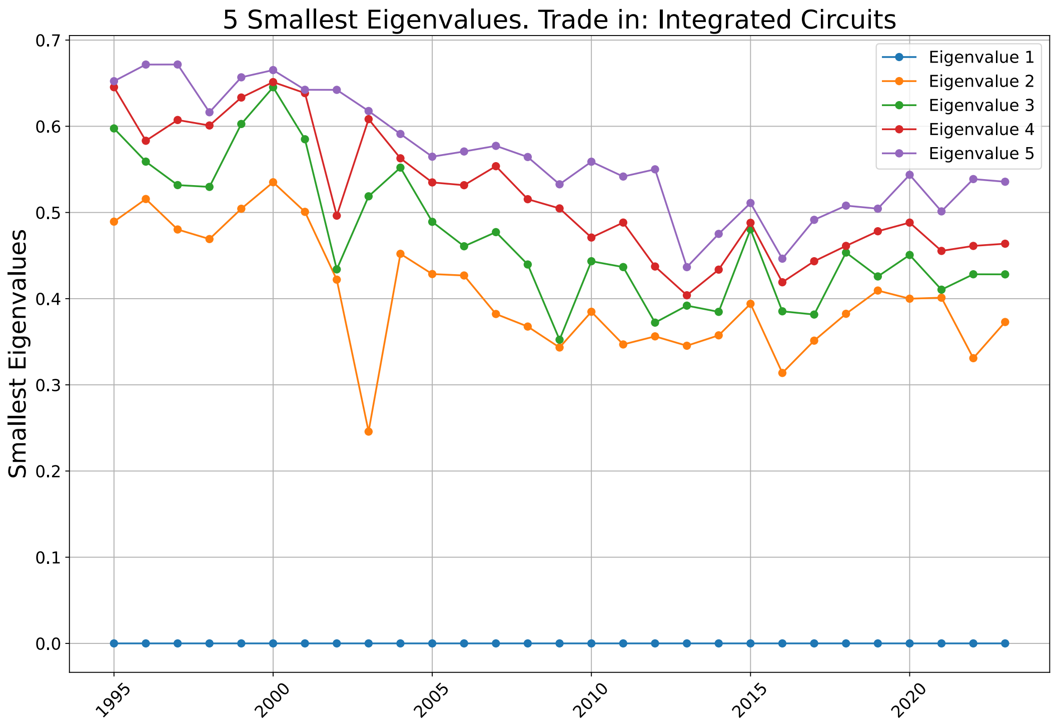

Figure 19: Five smallest eigenvalues of directed graph Laplacian $\mathcal{L}$ for trade in (8542). λ_2 in orange

.

indicating that trade in integrated circuits has become less connected over the last three decades. While we do not make causal claims, this trend coincides with the emergence of numerous new exporters in the sector, potentially leading to a more fragmented or modular supply chain. We also observe a sharp drop in λ_2 in 2003, coinciding with the outbreak of SARS in East Asia, which may have potentially strained supply chains at the time, contributing to this momentary decrease in network connectivity. Following the approach used in earlier sections, we also plot various metrics of network connectivity in Figure 42 where we plot λ_2, the approximated Φ_G, the real part of P's spectral, and the upper and lower bounds on Φ_G given by Cheeger's inequality. We observe general co-movement among these metrics across most years, with exceptions in 2018 and 2023. Understanding why the metrics move in different directions for just some years would require a deeper dive into the network topology. Finally, we run hierarchical clustering on the eigenvector corresponding to λ_2 for one year of network data in Figure 41 to show the partition of countries in the trade network.

5.2 Trade in Vehicles

Like above, we analyze trade in vehicles and vehicle parts by aggregating flows to the first two digits of HS codes, retaining all flows classified under code 87, which captures both trade in vehicle parts and final products. For each year of trade data, we compute the weighted export PageRank (PR) scores for all countries and plot the results in Figure 20. We also compute the difference between each country's PR score and its share of global vehicle exports, shown in Figure 43. In terms of centrality, Germany remains the dominant node in the global vehicle trade network throughout the entire sample period. We also observe a substantial rise in the centrality of China following its ascension into the WTO in 2001. Examining differences between PR and export share, countries like Germany and Japan are more central in the network than their trade share would suggest, and the US is significantly less central despite how much they export. This divergence may reflect the geographic concentration of US vehicle trade, which is heavily oriented toward regional partners such as Mexico and Canada, with relatively few vehicle exports leaving the North American region. Turning now to changes in network connectivity, we plot the change in the eigenvalues of

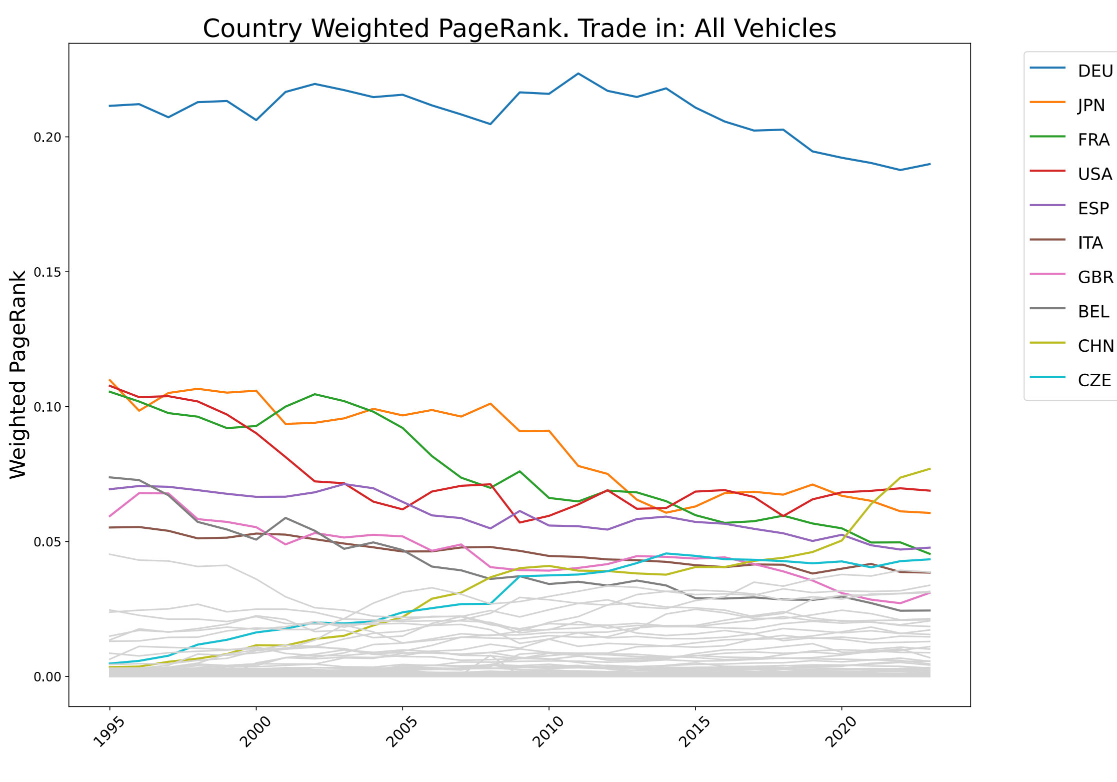

Figure 20: Weighted PageRank scores shown for countries in trade of vehicles (87). Transpose of the weight matrix used to identify exporters. $\alpha = 0.99$

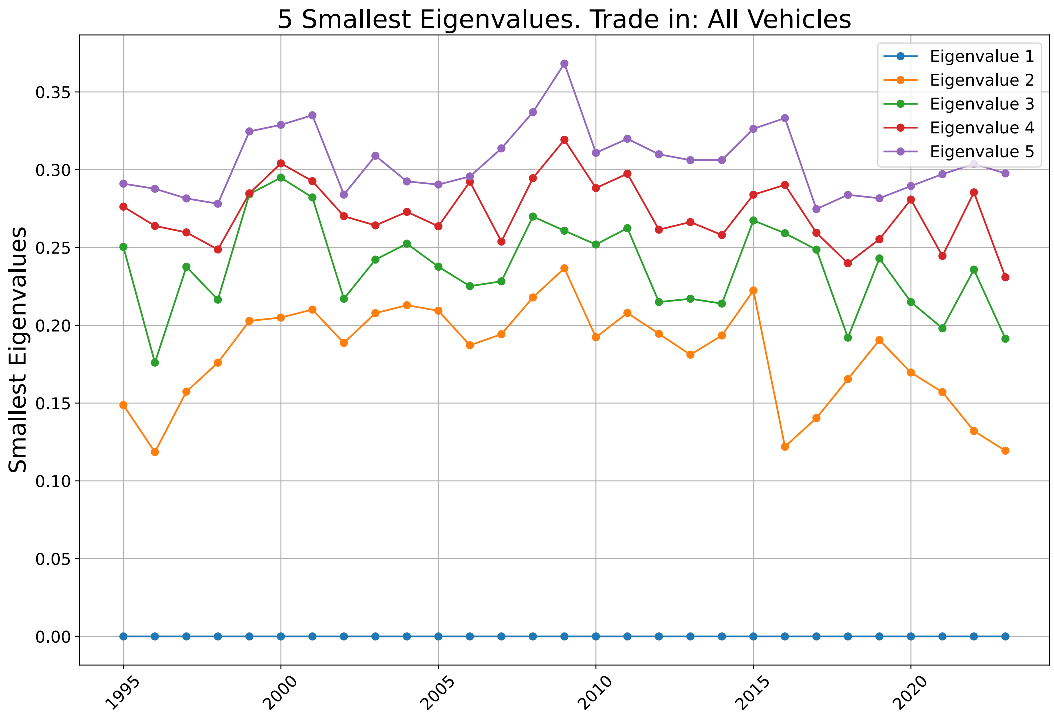

Figure 21: Five smallest eigenvalues of directed graph Laplacian $\mathcal{L}$ for trade in (87). λ_2 in orange .

the directed Laplacian of the vehicle trade network for each year of data in Figure 21.

Turning to network connectivity, we compute the directed Laplacian, $\mathcal{L}$, for the vehicle trade network in each year and plot its five smallest eigenvalues in Figure 21. In this chart, λ_2 shows an upward trend from 1995 to 2015, before dropping to its lowest value in 2016. Lastly, we apply spectral clustering to the eigenvector corresponding to λ_2 for the 2023 trade network as shown in Figure 44. The resulting partition highlights the regional fragmentation of the network, with North America forming one cluster, East Asia another, and Western and Northern Europe appearing as a distinct, tightly integrated bloc.

These analyses of the integrated circuit and vehicle trade networks serve as illustrative examples of how the combination of metrics outlined in this paper can be readily applied to the analysis of many types of trade networks.

6 Concluding Remarks

This study analyzed the evolution of the international trade network from 1950 to 2014 by applying tools from spectral graph theory and network centrality analysis. By constructing yearly weighted directed graphs representing world and intra-European trade, we analyzed the evolution in network connectivity using the eigenvalues of the directed graph Laplacians and graph conductance, and changes in country centrality as exporters using weighted PageRank. In addition to direct computation, we compared the observed network connectivity to simulated null models based on each network's existing weight matrix.

In the analysis of export centrality using weighted PageRank, we capture the decline of historical manufacturing hubs like the US and UK, alongside the ascent of China and Russia as more central exporters. Additionally, by comparing PageRank centrality with export shares, we identified Russia and the Netherlands as being more central than their trade flows would indicate as a result of the network structure of intra-European trade.

Using the second smallest eigenvalue, λ_2, of a network's directed Laplacian $\mathcal{L}$, we were unable to capture significant changes in the connectivity of the world network. However, we did identify a significant increase in connectivity of the intra-European network of trade following the reunification of Germany and the dissolution of the USSR.

Methodologically, this paper provides an overview of several metrics that, as far as we are aware, have not been previously applied to empirically analyze the connectivity trade networks, such as directed Laplacians and graph conductance. We hope that this thesis serves as a reference for those interested in applying these metrics to other economic networks. A natural extension of the work is an application of these methods to more granular network data, such as inter-industry inter-country flows or firm-to-firm flows across a variety of supply chains.

References

[Abbate et al., 2018] Abbate, A., De Benedictis, L., Fagiolo, G., and Tajoli, L. (2018). Distance-varying assortativity and clustering of the international trade networkADDENDUM. *Network Science*, 6(4):633–633.

[Aliniaeifard et al., 2021] Aliniaeifard, F., Wang, V., and van Willigenburg, S. (2021). Deletion-contraction for a unified Laplacian and applications. Version Number: 1.

[Baier et al., 2014] Baier, S. L., Bergstrand, J. H., and Feng, M. (2014). Economic integration agreements and the margins of international trade. *Journal of International Economics*, 93(2):339–350.

[Baldwin, 2006] Baldwin, R. E. (2006). The Euro's Trade Effects. *SSRN Electronic Journal*.

[Baltz and Kliemann, 2005] Baltz, A. and Kliemann, L. (2005). Spectral Analysis. In Hutchison, D., Kanade, T., Kittler, J., Kleinberg, J. M., Mattern, F., Mitchell, J. C., Naor, M., Nierstrasz, O., Pandu Rangan, C., Steffen, B., Sudan, M., Terzopoulos, D., Tygar, D., Vardi, M. Y., Weikum, G., Brandes, U., and Erlebach, T., editors, *Network Analysis*, volume 3418, pages 373–416. Springer Berlin Heidelberg, Berlin, Heidelberg. Series Title: Lecture Notes in Computer Science.

[Bapat, 2010] Bapat, R. B. (2010). *Graphs and matrices*. Universitext. Springer ; Hindustan Book Agency, London ; New York : New Delhi. OCLC: ocn455828013.

[Barabási and Pósfai, 2016] Barabási, A.-L. and Pósfai, M. (2016). *Network science*. Cambridge University Press, Cambridge, United Kingdom. OCLC: ocn910772793.

[Bhattacharya et al., 2007] Bhattacharya, K., Mukherjee, G., and Manna, S. S. (2007). The International Trade Network. In Chatterjee, A. and Chakrabarti, B. K., editors, *Econophysics of Markets and Business Networks*, pages 139–147. Springer Milan, Milano. Series Title: New Economic Windows.

[Brin and Page, 1998] Brin, S. and Page, L. (1998). The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Computer Networks*, 30:107–117.

[Butler, 2007] Butler, S. (2007). Interlacing for weighted graphs using the normalized Laplacian. *The Electronic Journal of Linear Algebra*, 16.

[Chaney, 2014] Chaney, T. (2014). The Network Structure of International Trade. *The American Economic Review*, 104(11):3600–3634. Publisher: American Economic Association.

[Chaney, 2016] Chaney, T. (2016). Networks in International Trade. In Bramoullé, Y., Galeotti, A., and Rogers, B. W., editors, *The Oxford Handbook of the Economics of Networks*, pages 753–775. Oxford University Press.

[Chatterjee, 2023] Chatterjee, S. (2023). Spectral gap of nonreversible Markov chains. Version Number: 3.

[Chen et al., 2004] Chen, G., Davis, G., Hall, F., Li, Z., Patel, K., and Stewart, M. (2004). An Interlacing Result on Normalized Laplacians. *SIAM Journal on Discrete Mathematics*, 18(2):353–361. Publisher: Society for Industrial and Applied Mathematics.

[Chung, 1997] Chung, F. R. (1997). *Spectral graph theory: CBMS Conference on Recent Advances in Spectral Graph Theory, held at ... Fresno, June 6 - 10, 1994*. Number 92 in Regional conference series in mathematics. American Mathematical Society, Providence, RI, reprint edition. Meeting Name: Conference on Recent Advances in Spectral Graph Theory.

[Chung, 2005] Chung, F. R. (2005). Laplacians and the Cheeger Inequality for Directed Graphs.

[De Benedictis et al., 2014] De Benedictis, L., Nenci, S., Santoni, G., Tajoli, L., and Vicarelli, C. (2014). Network Analysis of World Trade Using the BACI-CEPII Dataset. *SSRN Electronic Journal*.

[De Clerck et al., 2024] De Clerck, B., Van Utterbeeck, F., and Rocha, L. E. C. (2024). Comparative analysis of graph randomization: Tools,methods, pitfalls, and best practices. Version Number: 1.

[Deguchi et al., 2014] Deguchi, T., Takahashi, K., Takayasu, H., and Takayasu, M. (2014). Hubs and Authorities in the World Trade Network Using a Weighted HITS Algorithm. *PLoS ONE*, 9(7):e100338.

[Elliott and Golub, 2022] Elliott, M. and Golub, B. (2022). Networks and Economic Fragility. *Annual Review of Economics*, 14(1):665–696.

[Erdös and Renyi, 1959] Erdös, P. and Renyi, A. (1959). On Random Graphs I. *Publicationes Mathematicate*, (6):290–297.

[Fagiolo, 2006] Fagiolo, G. (2006). Directed or Undirected? A New Index to Check for Directionality of Relations in Socio-Economic Networks. Publisher: arXiv Version Number: 3.

[Fagiolo, 2010] Fagiolo, G. (2010). The international-trade network: gravity equations and topological properties. *Journal of Economic Interaction and Coordination*, 5(1):1–25.

[Fagiolo, 2016] Fagiolo, G. (2016). *The International Trade Network*, volume 1. Oxford University Press.

[Fagiolo et al., 2010] Fagiolo, G., Reyes, J., and Schiavo, S. (2010). The evolution of the world trade web: a weighted-network analysis. *Journal of Evolutionary Economics*, 20(4):479–514.

[Fiedler, 1973] Fiedler, M. (1973). Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305. Publisher: Institute of Mathematics, Academy of Sciences of the Czech Republic.

[Fisk, 2005] Fisk, S. (2005). A very short proof of Cauchy's interlace theorem for eigenvalues of Hermitian matrices. Publisher: arXiv Version Number: 1.

[Fouquin and Hugot, 2016] Fouquin, M. and Hugot, J. (2016). Two Centuries of Bilateral Trade and Gravity Data: 1827-2014.

[Gaulier and Zignago, 2025] Gaulier, G. and Zignago, S. (2025). BACI: International Trade Database at the Product-Level. The 1994-2007 Version. Working Papers 2010-23, CEPII.

[Gleich, 2006] Gleich, D. F. (2006). Hierarchical Directed Spectral Graph Partitioning MS&E 337 - Information Networks.

[Golub and Van Loan, 1983] Golub, G. H. and Van Loan, C. F. (1983). *Matrix computations*. Johns Hopkins studies in the mathematical sciences. The Johns Hopkins University Press, Baltimore, fourth edition edition. OCLC: ocn824733531.

[Hoang et al., 2023] Hoang, V. P., Piccardi, C., and Tajoli, L. (2023). Reshaping the structure of the World Trade Network: a pivotal role for China? *Applied Network Science*, 8(1):35.

[Ioannides, 1996] Ioannides, Y. M. (1996). Introduction. *Regional Science and Urban Economics*, 26(6):561–563.

[Ioannides, 2015] Ioannides, Y. M. (2015). Random Graphs and Social Networks: An Economics Perspective.

[Jackson, 2011] Jackson, M. O. (2011). *Social and economic networks*. Princeton University Press, Princeton, N.J. Woodstock.

[Kabir et al., 2017] Kabir, M., Salim, R., and Al-Mawali, N. (2017). The gravity model and trade flows: Recent developments in econometric modeling and empirical evidence. *Economic Analysis and Policy*, 56:60–71.

[Kim and Yun, 2022] Kim, S. and Yun, J. (2022). Analysis of risk propagation using the world trade network. *Journal of the Korean Physical Society*, 81(7):697–706.

[Kivela et al., 2014] Kivela, M., Arenas, A., Barthelemy, M., Gleeson, J. P., Moreno, Y., and Porter, M. A. (2014). Multilayer networks. *Journal of Complex Networks*, 2(3):203–271.

[Kleinberg, 1999] Kleinberg, J. M. (1999). Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46(5):604–632.

[Korniyenko et al., 2017] Korniyenko, Y., Pinat, M., and Dew, B. (2017). *Assessing the fragility of global trade: the impact of localized supply shocks using network analysis*. International Monetary Fund, Washington, District of Columbia. OCLC: 982375430.

[Krugman, 1979] Krugman, P. R. (1979). Increasing returns, monopolistic competition, and international trade. *Journal of International Economics*, 9(4):469–479.

[Laenen and Sun, 2020] Laenen, S. and Sun, H. (2020). Higher-Order Spectral Clustering of Directed Graphs. Publisher: [object Object] Version Number: 1.

[Levin et al., 2009] Levin, D. A., Peres, Y., and Wilmer, E. L. (2009). *Markov chains and mixing times*. American Mathematical Society, Providence, Rhode Island.

[Pao et al., 2011] Pao, H., Coppersmith, G. A., and Priebe, C. E. (2011). Statistical Inference on Random Graphs: Comparative Power Analyses via Monte Carlo. *Journal of Computational and Graphical Statistics*, 20(2):395–416.

[Sajedianfard et al., 2021] Sajedianfard, N., Hadian, E., Samadi, A. H., Dehghan Shabani, Z., Sarkar, S., and Robinson, P. A. (2021). Quantitative analysis of trade networks: data and robustness. *Applied Network Science*, 6(1):46.

[Seabrook and Wiskott, 2023] Seabrook, E. and Wiskott, L. (2023). A Tutorial on the Spectral Theory of Markov Chains. *Neural Computation*, 35(11):1713–1796.

[Setayesh et al., 2022] Setayesh, A., Sourati Hassan Zadeh, Z., and Bahrak, B. (2022). Analysis of the global trade network using exponential random graph models. *Applied Network Science*, 7(1):38.

[Shi and Malik, 2000] Shi, J. and Malik, J. (2000). Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.

[Sorokin, 2022] Sorokin, A. (2022). The Soviet Unions Economic Relations with Austria and the Federal Republic of Germany: Political Factors and the Art of Diplomacy (1955-1964). *Quaestio Rossica*, 10(5).

[Spielman,] Spielman, D. *Spectral and Algebraic Graph Theory Incomplete Draft*.

[Squartini et al., 2013] Squartini, T., van Lelyveld, I., and Garlaschelli, D. (2013). Early-warning signals of topological collapse in interbank networks. Publisher: arXiv Version Number: 4.

[Taylor et al., 2015] Taylor, D., Myers, S. A., Clauset, A., Porter, M. A., and Mucha, P. J. (2015). Eigenvector-Based Centrality Measures for Temporal Networks. Version Number: 3.

[Tinbergen, 1962] Tinbergen, J. (1962). *Shaping the World Economy: Suggestions for an International Economic Policy*. Twentieth Century Fund, New York.

[United States. Department of Transportation. Bureau of Transportation Statistics, 1993] United States. Department of Transportation. Bureau of Transportation Statistics (1993). Commodity Flow Survey (CFS) 1993 [Supporting Datasets].

[United States. Department of Transportation. Bureau of Transportation Statistics, 1997] United States. Department of Transportation. Bureau of Transportation Statistics (1997). Commodity Flow Survey (CFS) 1997 [Supporting Datasets].

[United States. Department of Transportation. Bureau of Transportation Statistics, 2002] United States. Department of Transportation. Bureau of Transportation Statistics (2002). Commodity Flow Survey 2002 [supporting datasets].

[United States. Department of Transportation. Bureau of Transportation Statistics, 2007] United States. Department of Transportation. Bureau of Transportation Statistics (2007). Commodity Flow Survey 2007 [supporting datasets].

[United States. Department of Transportation. Bureau of Transportation Statistics, 2012] United States. Department of Transportation. Bureau of Transportation Statistics (2012). Commodity Flow Survey 2012 [supporting datasets].

[United States. Department of Transportation. Bureau of Transportation Statistics, 2017] United States. Department of Transportation. Bureau of Transportation Statistics (2017). Commodity Flow Survey 2017 [supporting datasets].

[Von Luxburg, 2007] Von Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416.

[Ward, 1963] Ward, J. H. (1963). Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association*, 58(301):236–244.

[Watts and Strogatz, 1998] Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of small-world networks. *Nature*, 393(6684):440–442.

[Williamson, 2016] Williamson, D. P. (2016). 6334 Spectral Graph Theory Lecture 4 [Lecture Notes].

[Wyss-Gallifent, 2019] Wyss-Gallifent, J. (2019). *Applications of Linear Algebra*.

[Zhang et al., 2022] Zhang, P., Wang, T., and Yan, J. (2022). PageRank centrality and algorithms for weighted, directed networks. *Physica A: Statistical Mechanics and its Applications*, 586:126438.

7 Appendix

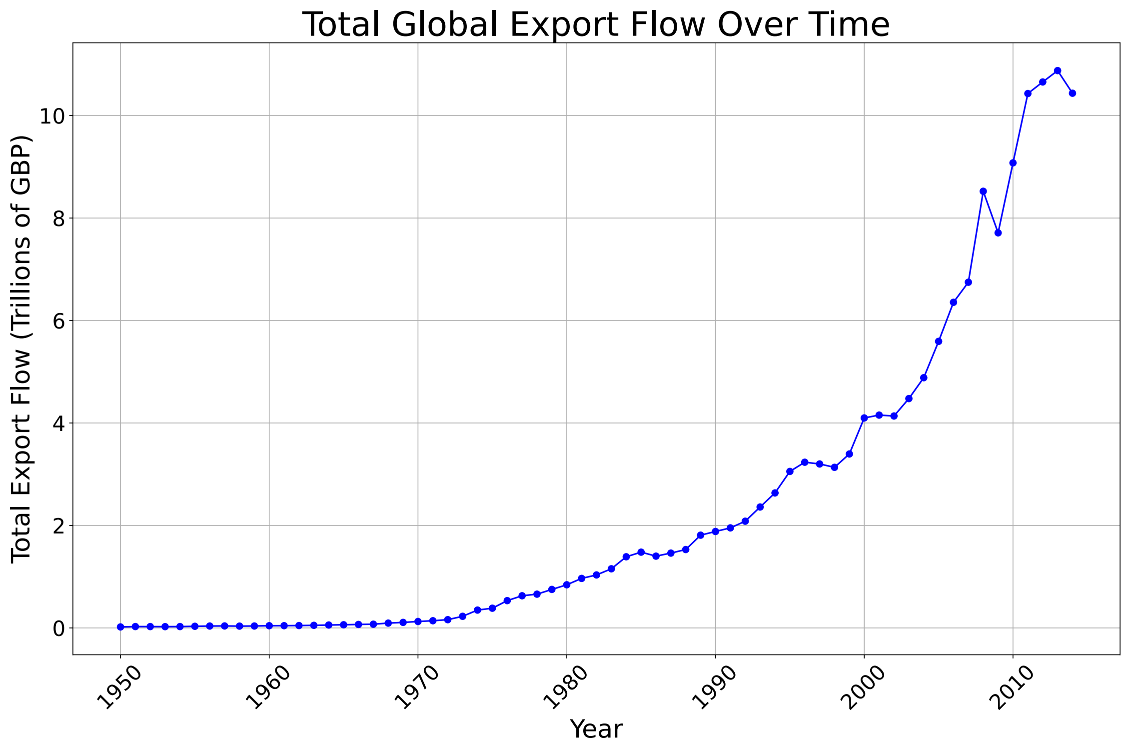

Figure 22: Total Export Flow by Year in Trillions of GBP

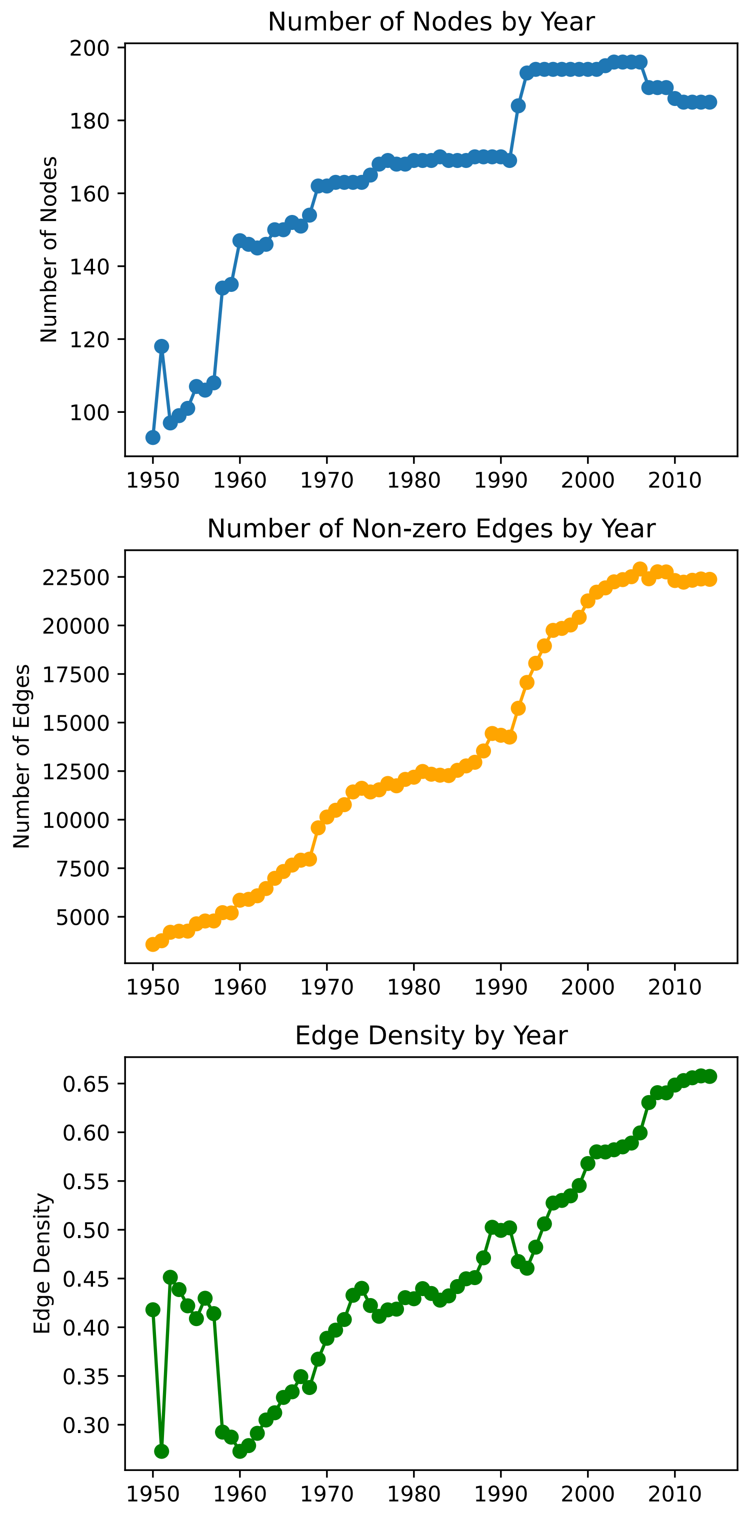

Figure 23: Summary Stats of W_t for years 1950-2014. Edge density is defined as $\frac{edges}{N*(N-1)}$. (Bhattacharya, 2007) used to select relevant summary statistics.

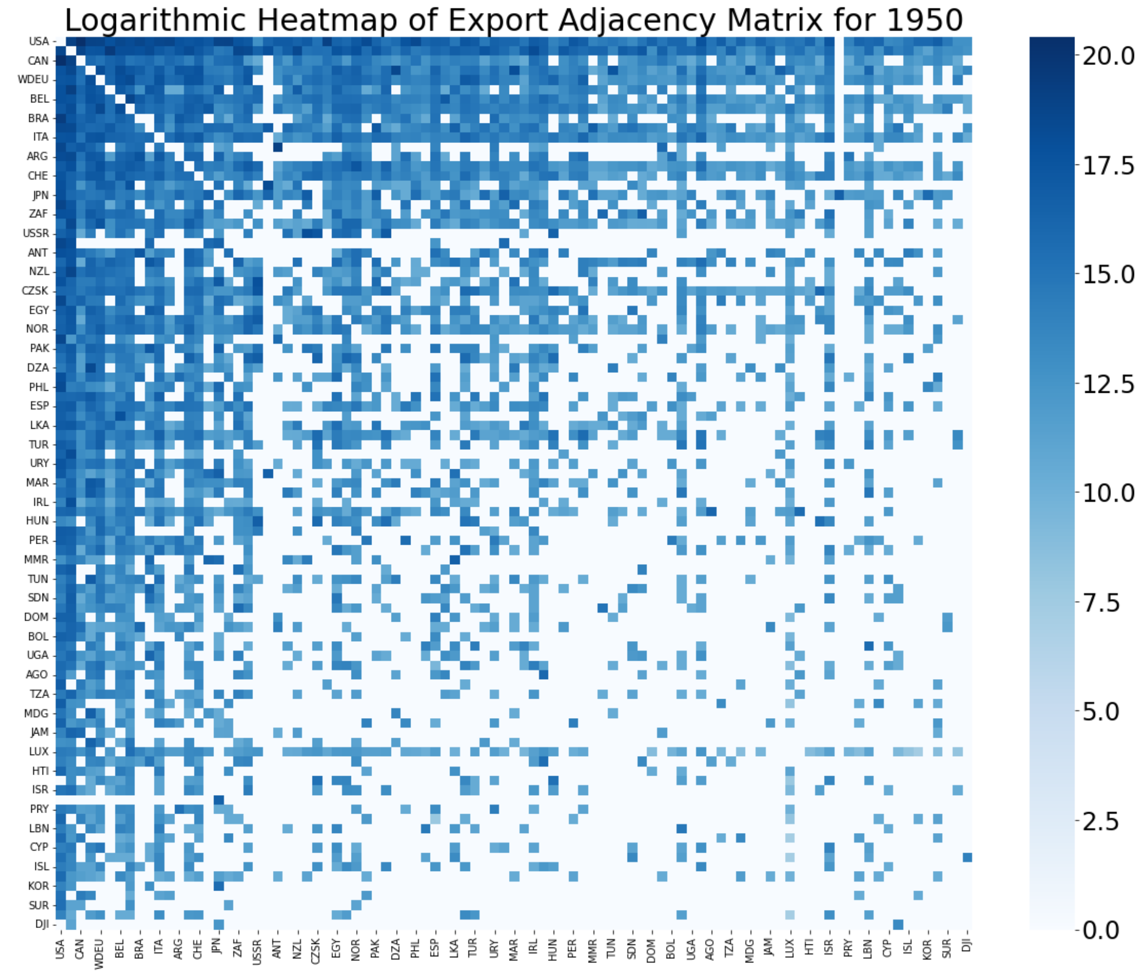

Figure 24: Logarithmic heat-map of W_t for year 1950. Rows and columns are ordered by the sum of a country's total exports to all countries.

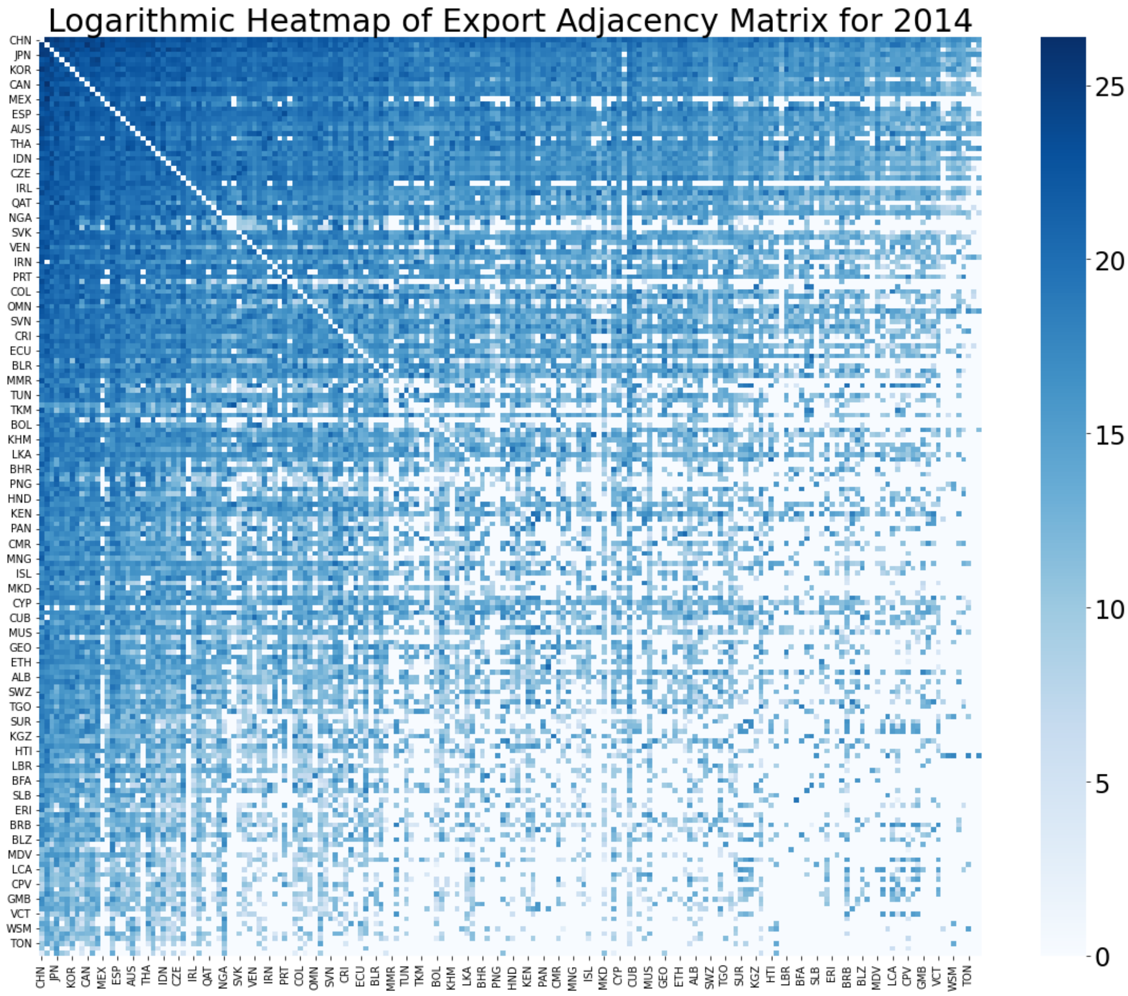

Figure 25: Logarithmic heat-map of W_t for year 2014. Rows and columns are ordered by the sum of a country's total exports to all countries.

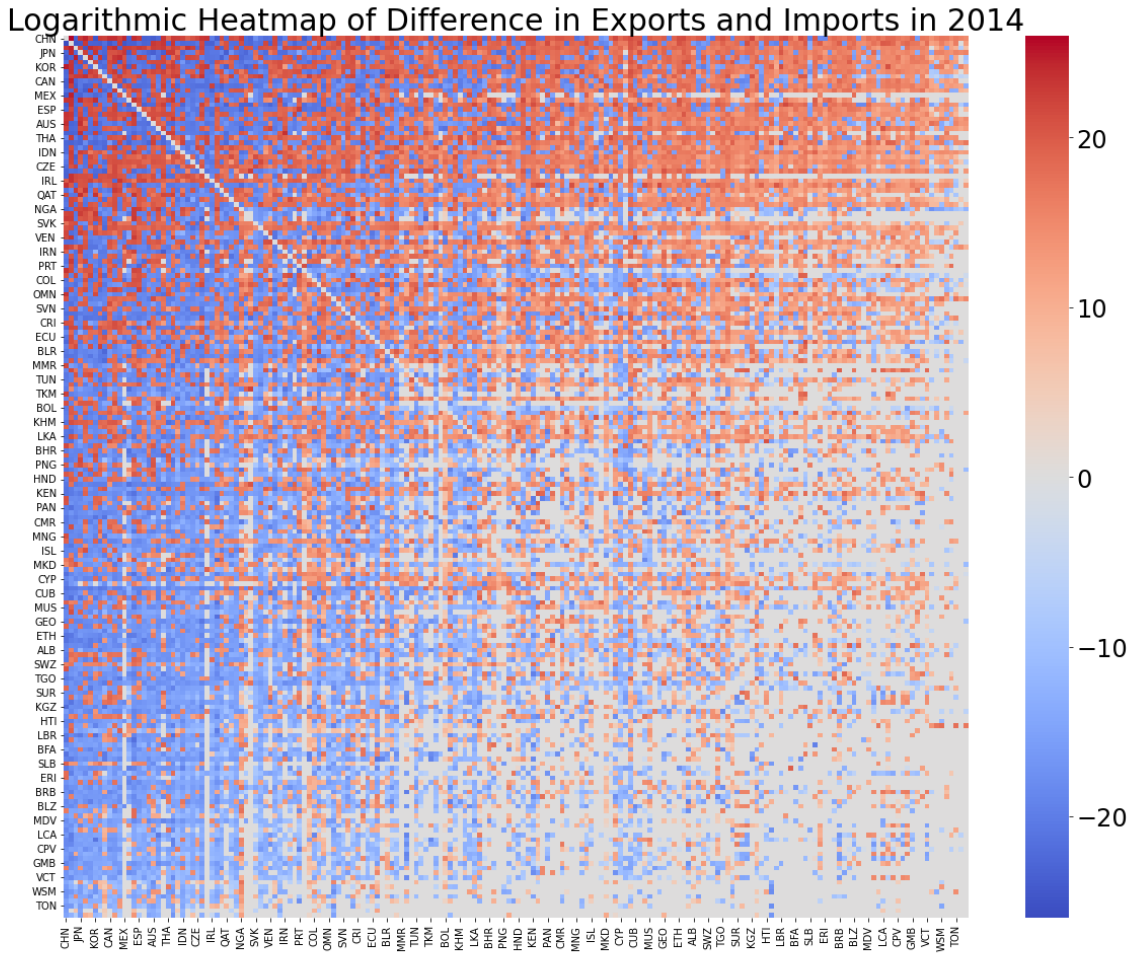

Figure 26: Logarithmic difference of exports and imports for the year 2014. Positive values represent trade surplus from country i to country j.

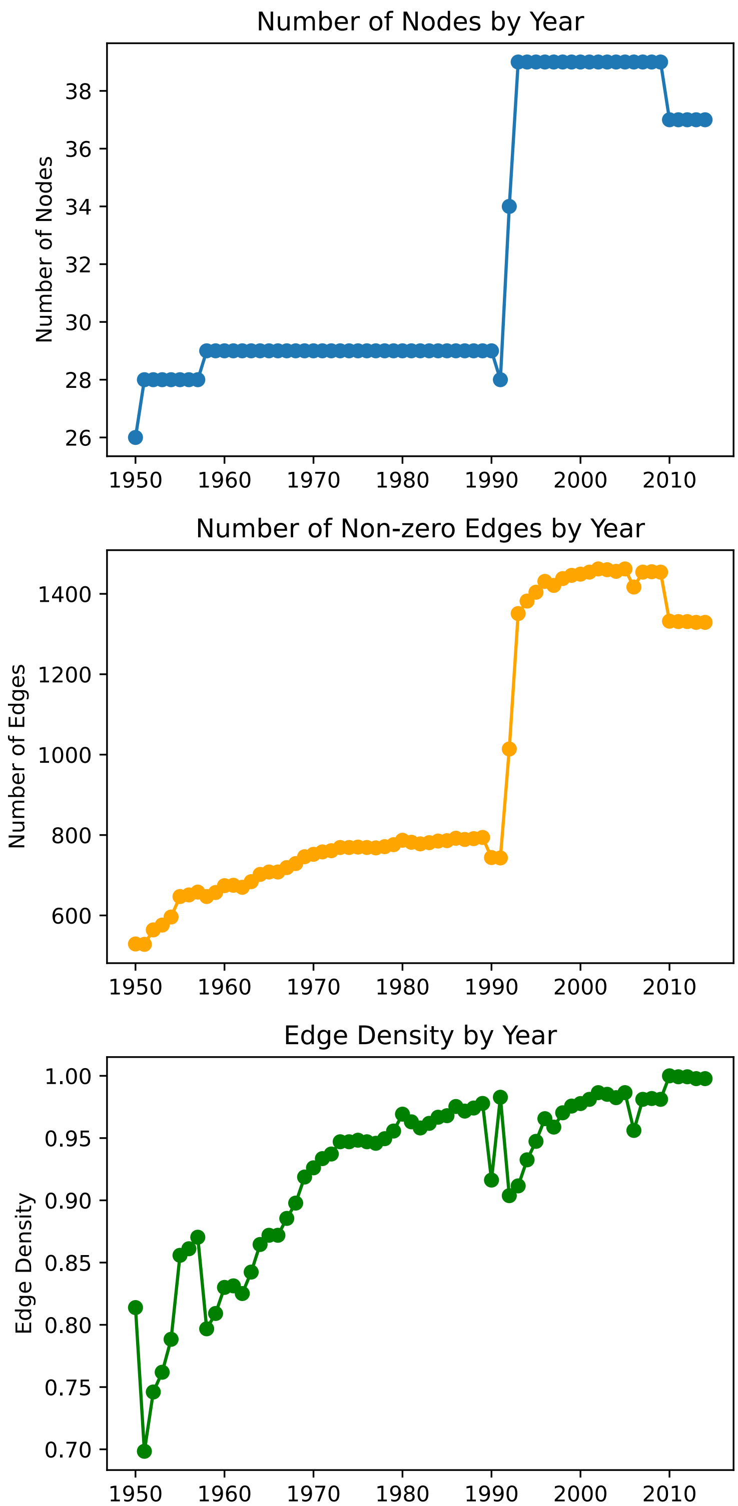

Figure 27: Intra-European trade network: Summary Stats of W_t for years 1950-2014. Edge density is defined as $\frac{edges}{N*(N-1)}$.

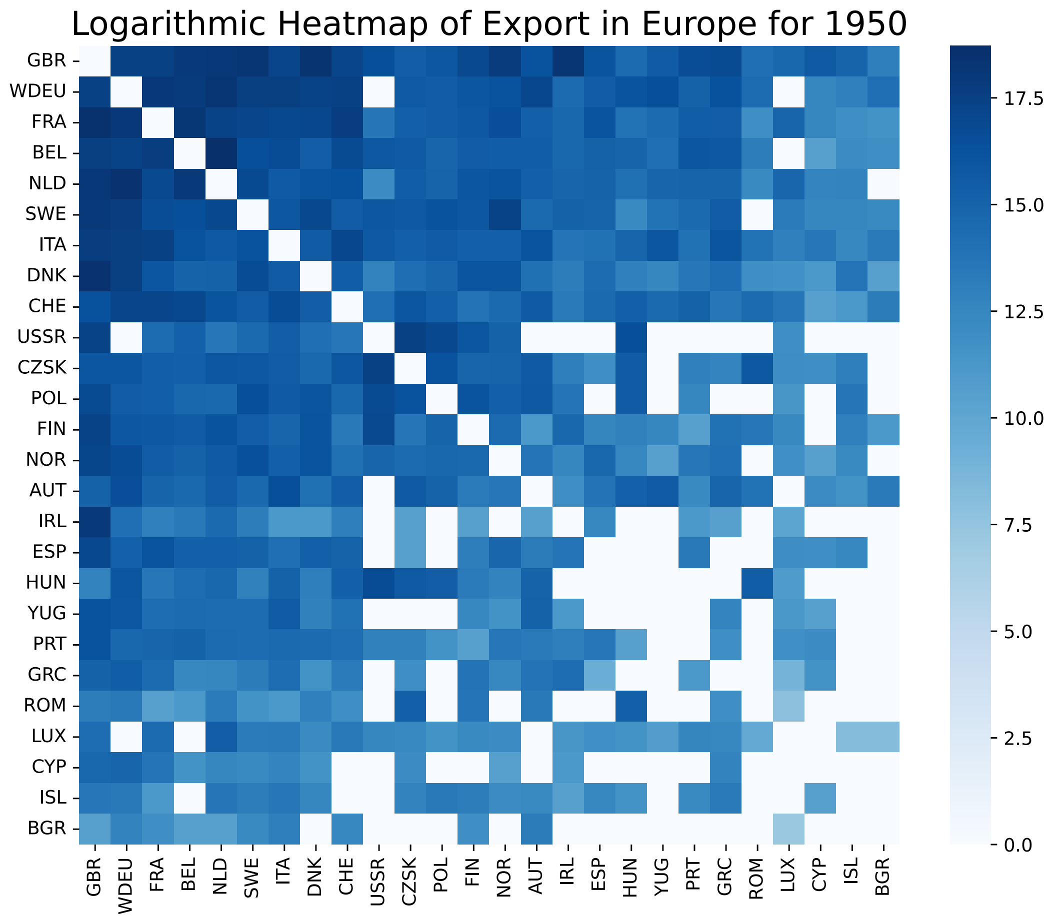

Figure 28: Logarithmic heat-map of intra-Europe trade for year 1950. Rows and columns are ordered by the sum of a country's total exports to all European countries.

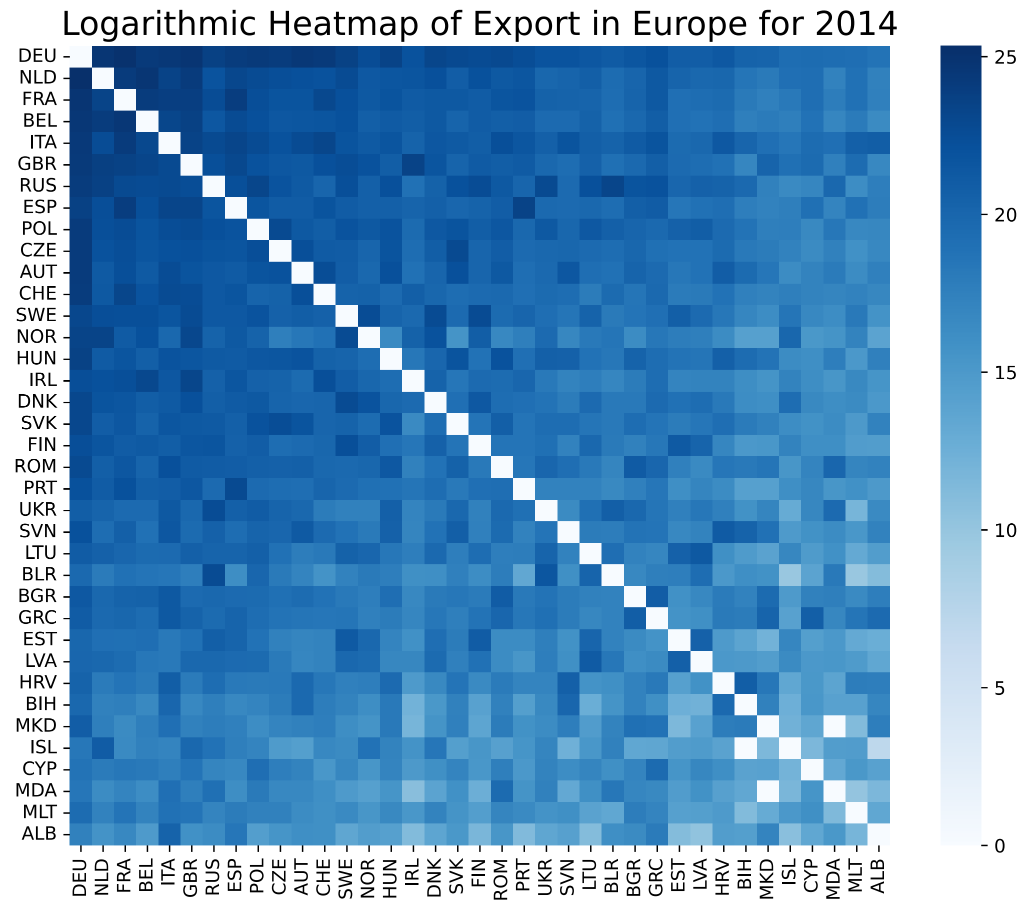

Figure 29: Logarithmic heat-map of intra-Europe trade for year 2014. Rows and columns are ordered by the sum of a country's total exports to all European countries.

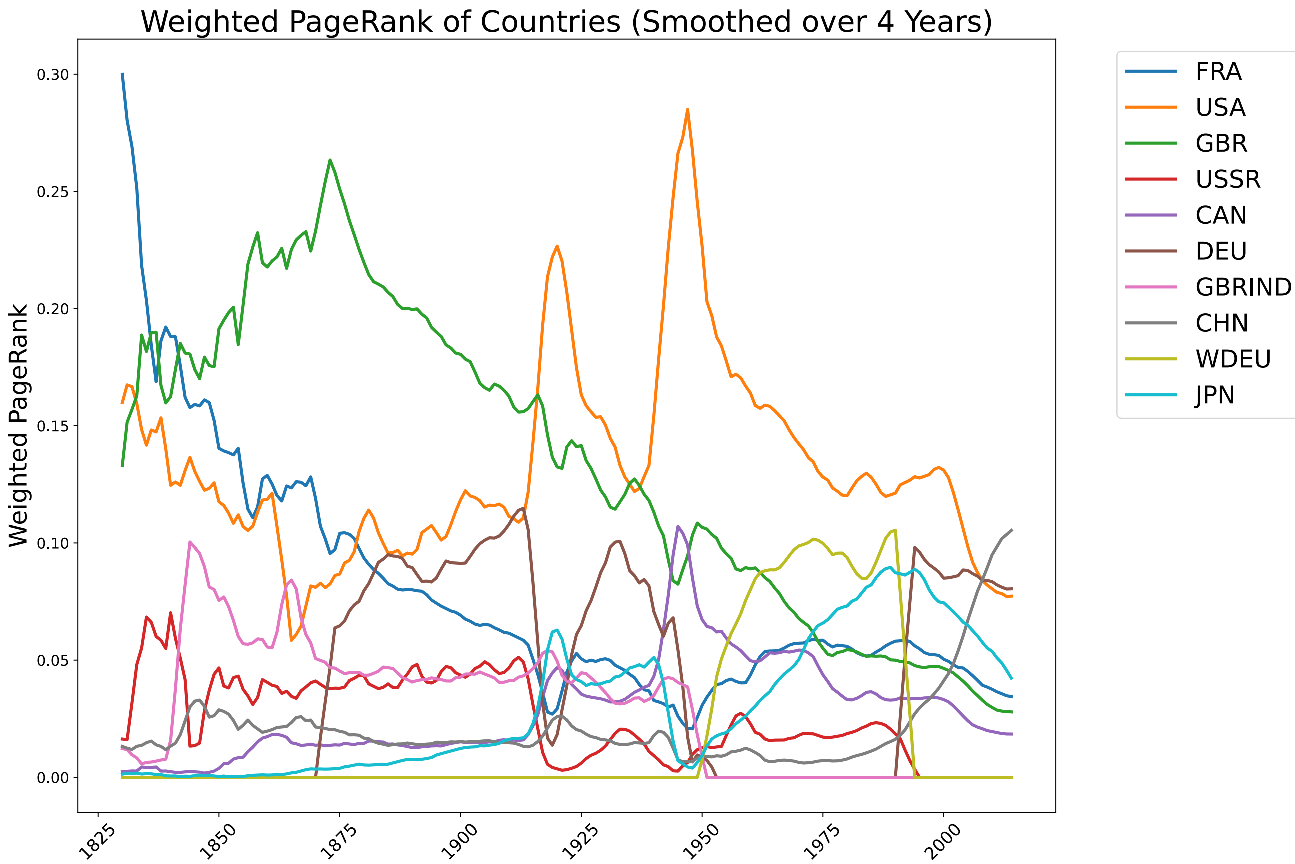

Figure 30: Weighted PageRank scores shown for the top 10 countries of the global trade network for years 1827-2014. PageRanks are computed using the transpose of the network's weight matrix. $\alpha = 0.99$

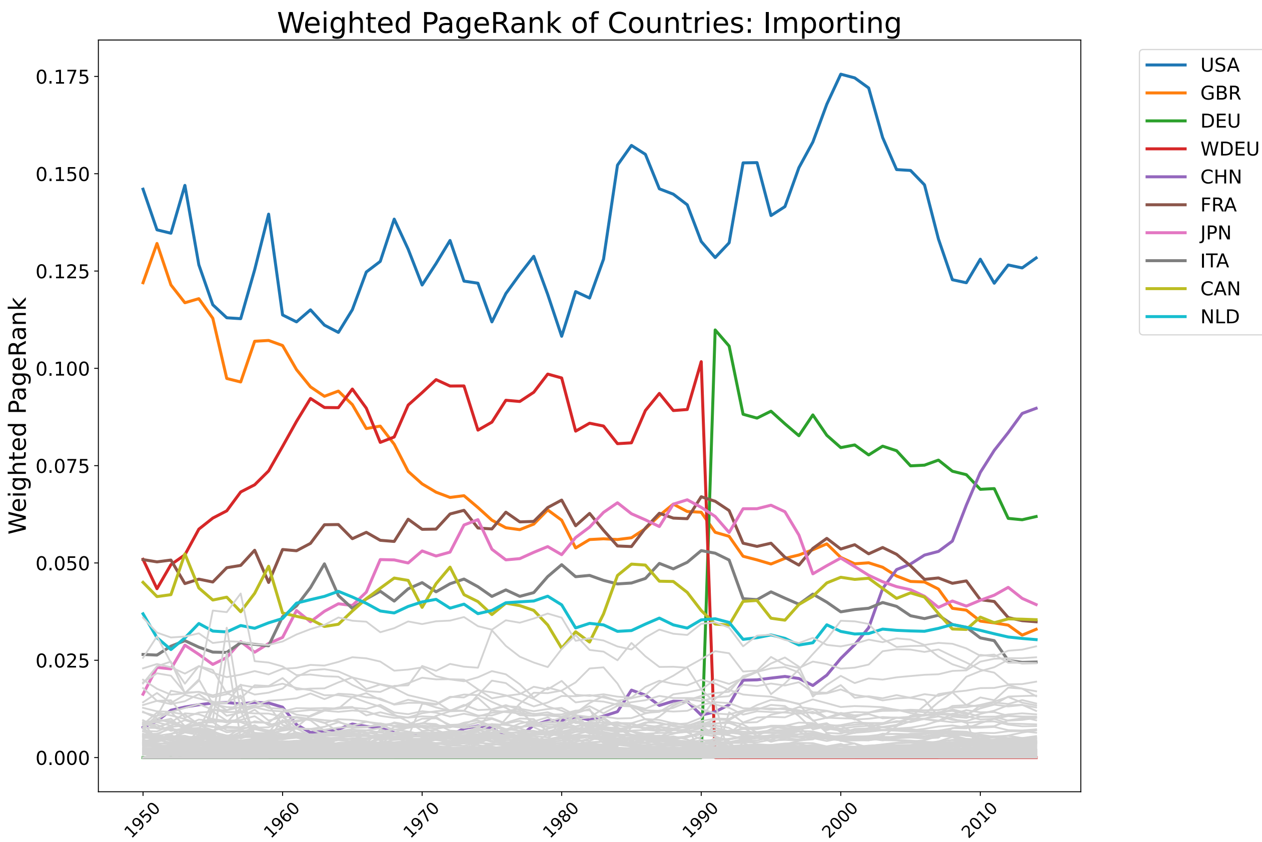

Figure 31: Weighted PageRank of countries in terms of imports with $\alpha = 0.95$. Countries with the top 10 average PageRank values are labeled.

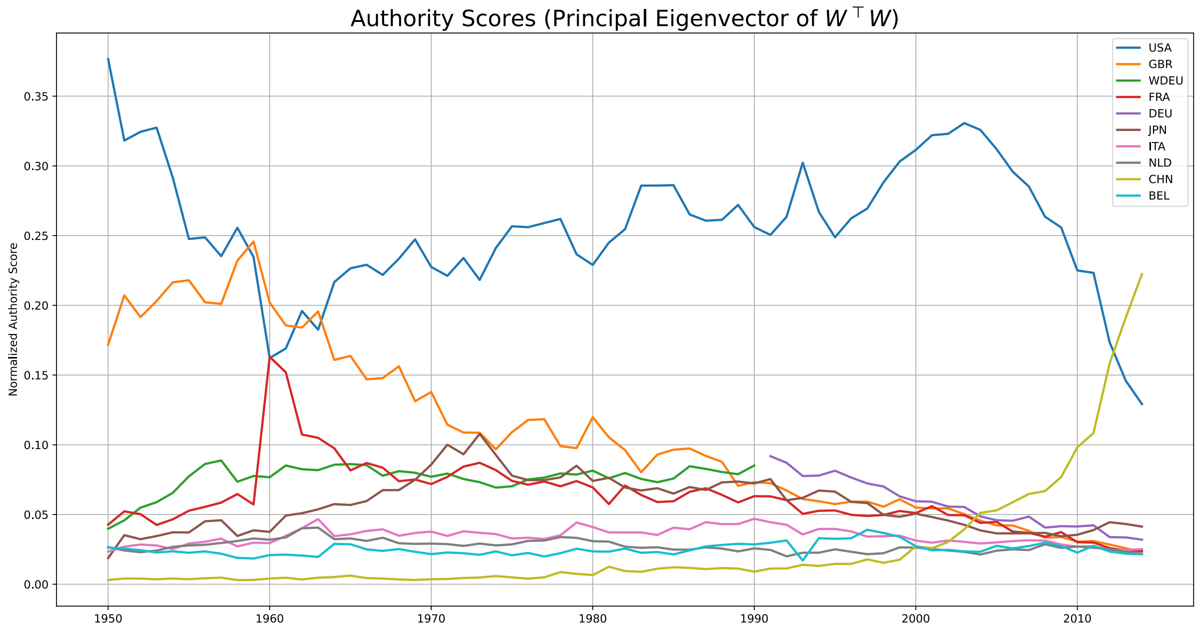

Figure 32: World trade network: Normalized authority scores computed from the principal eigenvector of $\hat{W}^\top \hat{W}$.

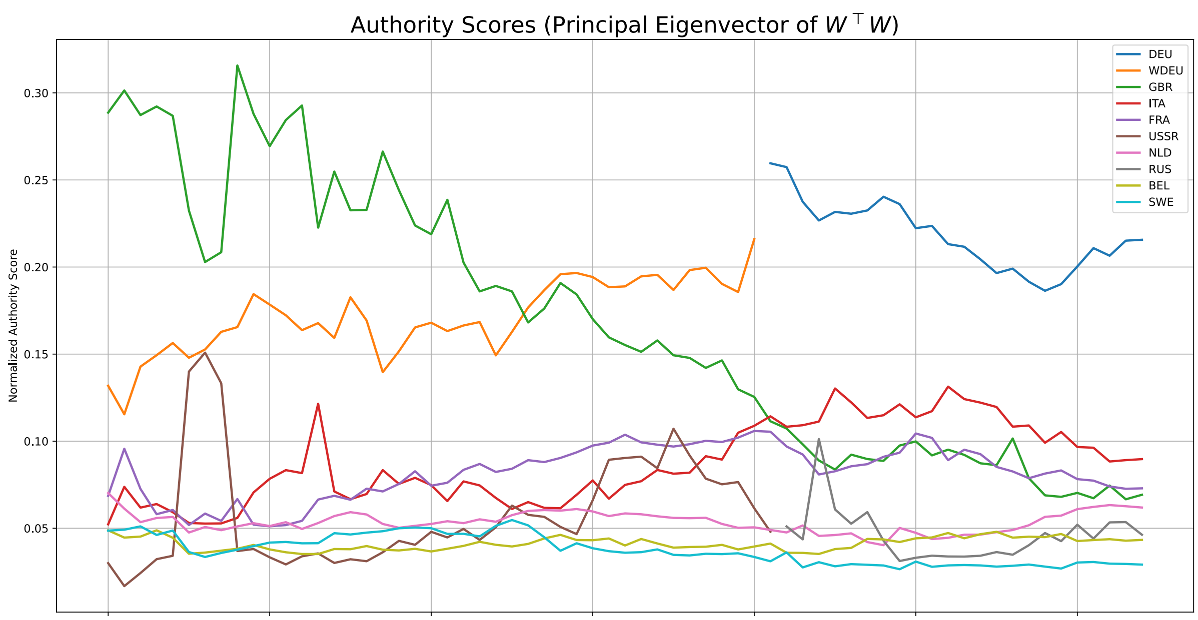

Figure 33: Intra-European trade network: Normalized authority scores computed from the principal eigenvector of $\hat{W}^\top \hat{W}$.

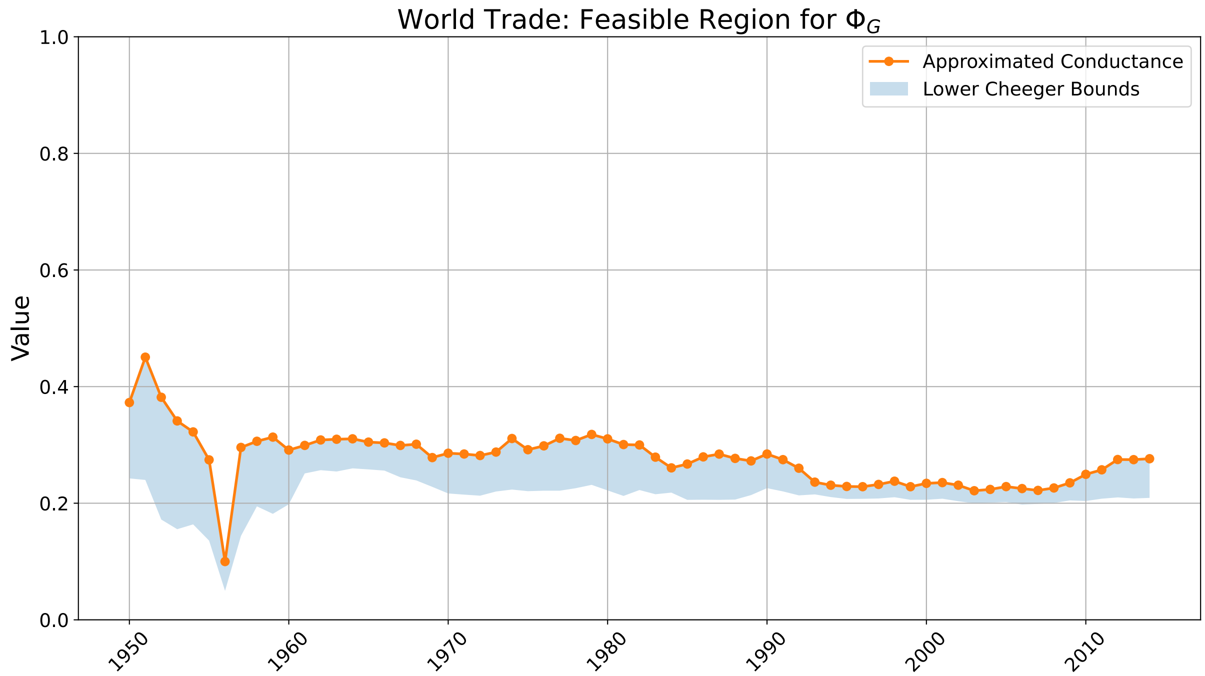

Figure 34: World trade: Feasible region for Φ_G using approximated value and lower bound of Cheeger inequality.

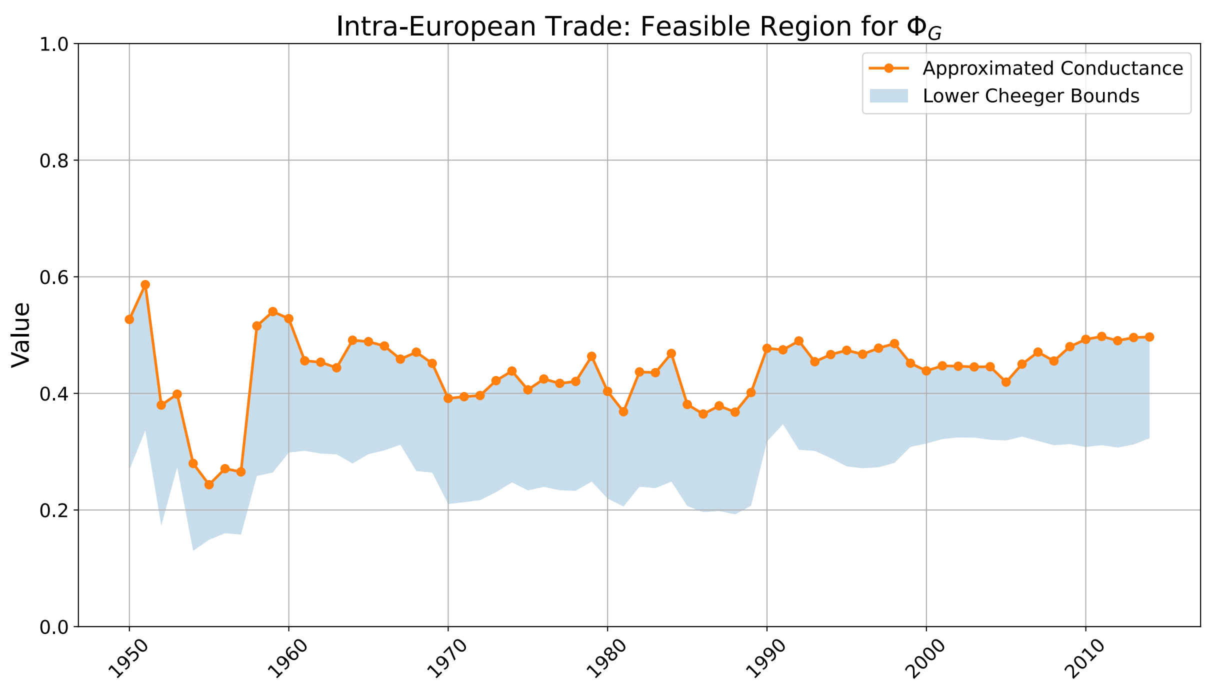

Figure 35: Intra-European trade: Feasible region for Φ_G using approximated value and lower bound of Cheeger inequality.

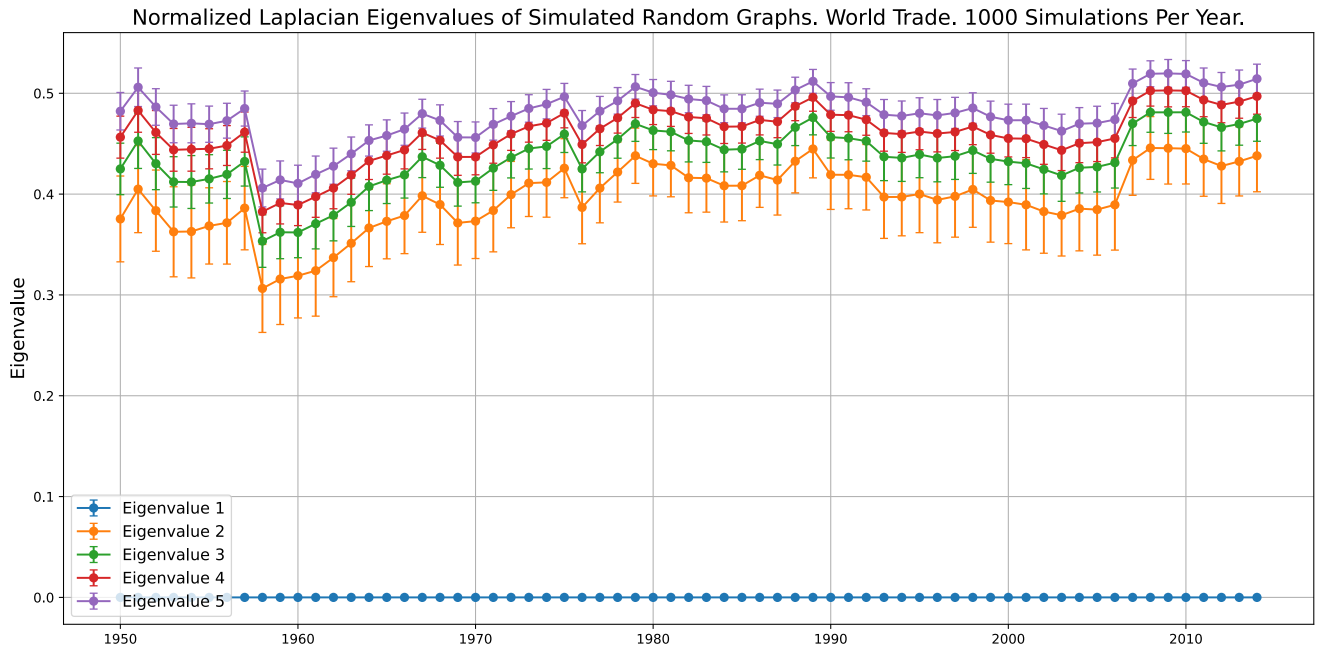

Figure 36: World Trade Network: Mean value of five smallest eigenvalues of the directed Laplacian for 1000 replications of the random graph $\mathcal{G}$ with parameter values from the full trade network.

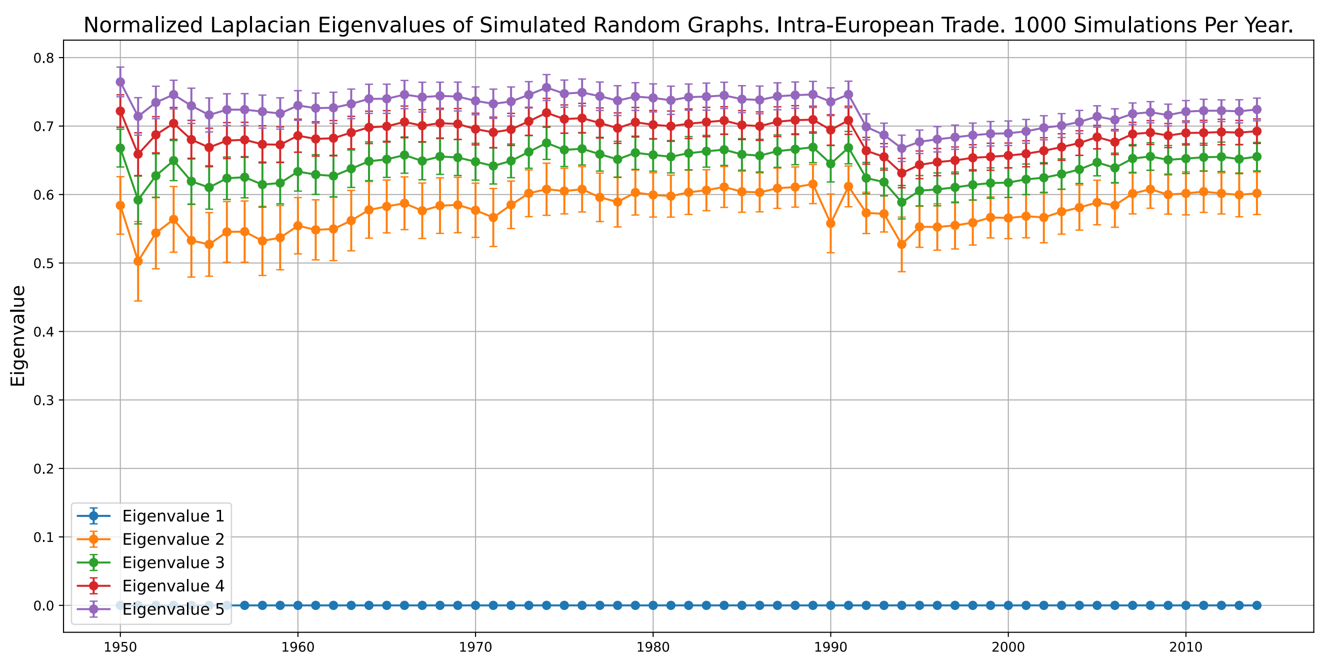

Figure 37: Intra-European Trade Network: Mean value of five smallest eigenvalues of the directed Laplacian for 1000 replications of the random graph $\mathcal{G}$ with parameter values from the intra-European trade network.

Clustered Network of Intra-European Trade in 1991 (Two Largest Export Destinations).

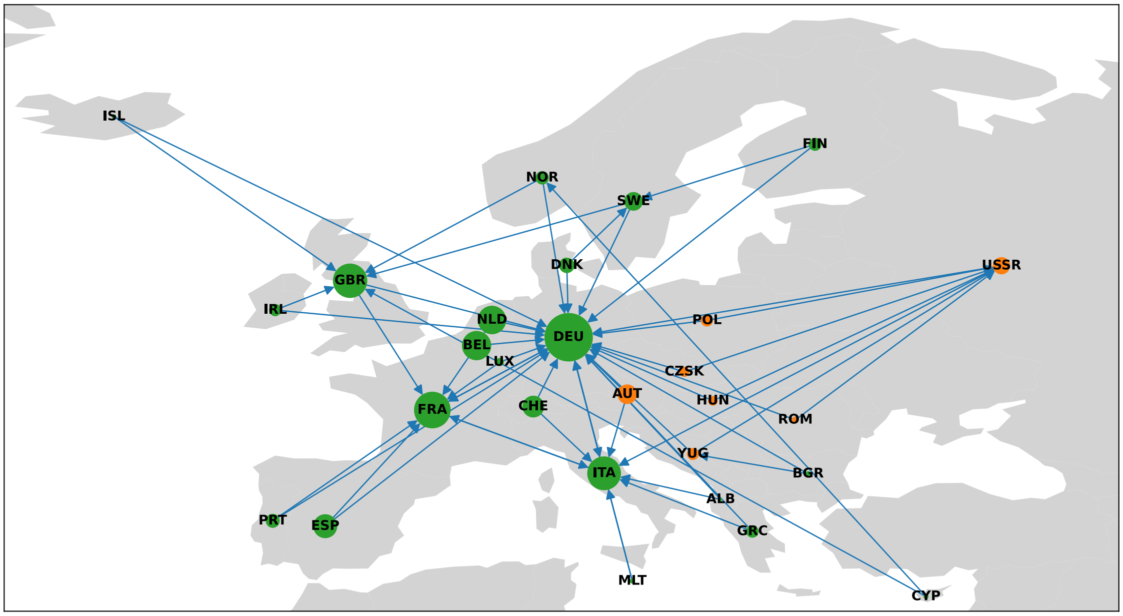

Figure 38: Directed edges are drawn for each country's two largest export destinations. Country sizes are based on their in-degree, colors based on the first hierarchical clustering partition.

Clustered Network of Intra-European Trade in 1992 (Two Largest Export Destinations).

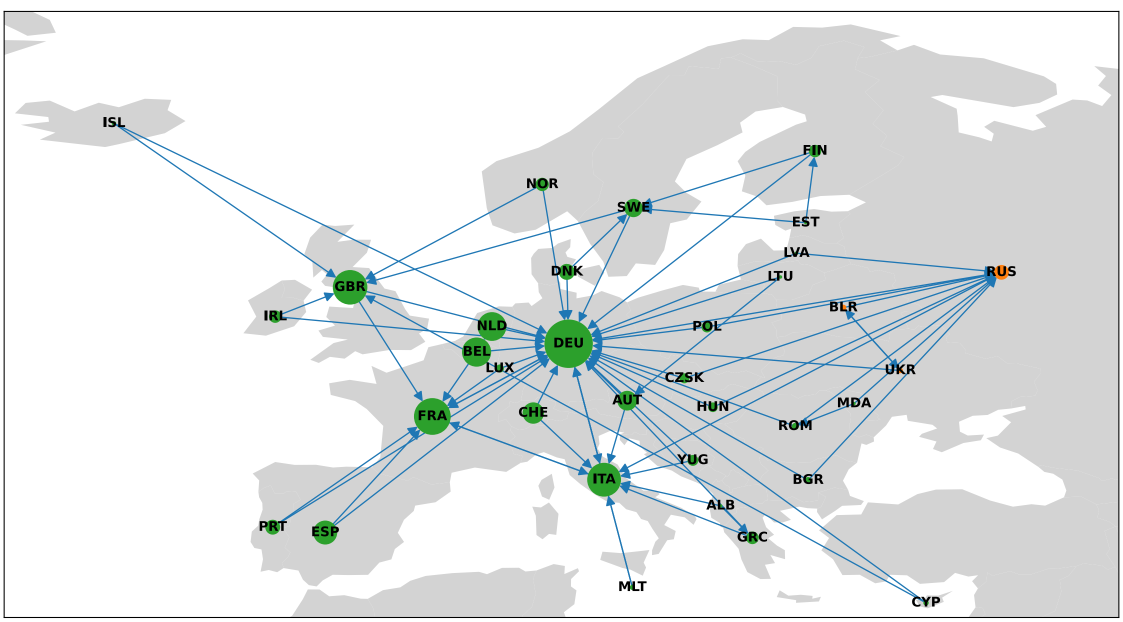

Figure 39: Directed edges are drawn for each country's two largest export destinations. Country sizes are based on their in-degree, colors based on the first hierarchical clustering partition.

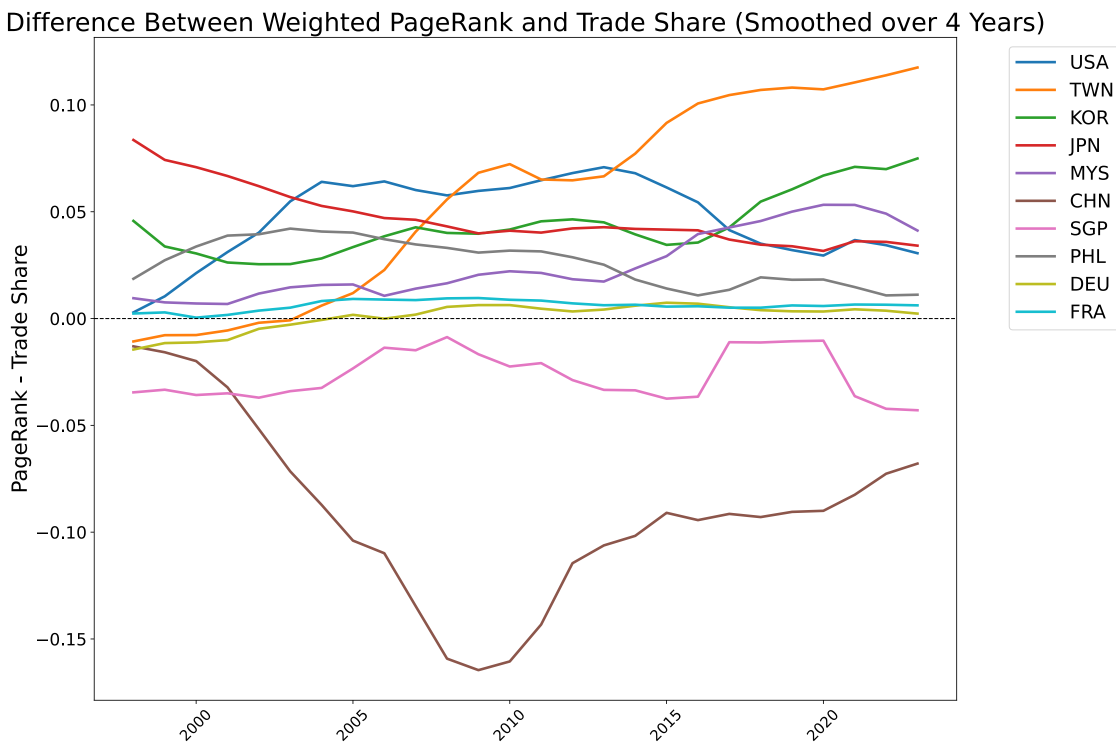

Figure 40: Weighted PageRank - Export Share for top 10 countries of the integrated circuits trade network. PageRanks are computed using the transpose of the network's weight matrix. $\alpha = 0.99$

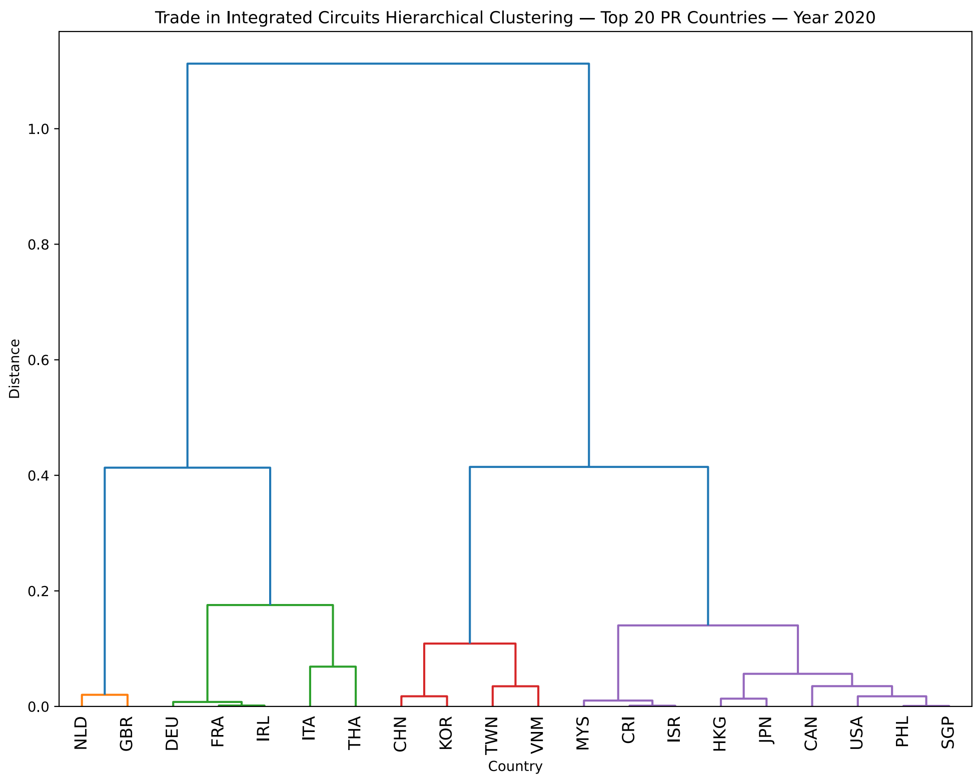

Figure 41: Hierarchical Clustering in 2020 using the eigenvector corresponding to λ_2 for the 20 countries with the highest average weighted PageRanks.

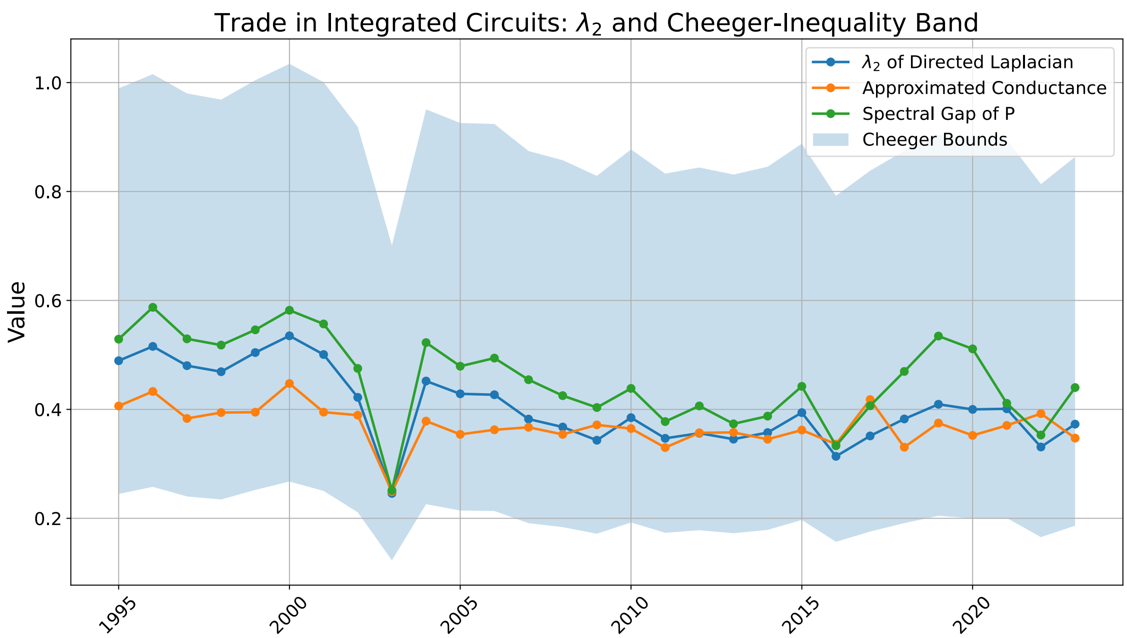

Figure 42: Integrated Circuits: λ_2 of directed Laplacian, approximated Φ_G, real part of P's spectral gap, and Cheeger's inequality bands.

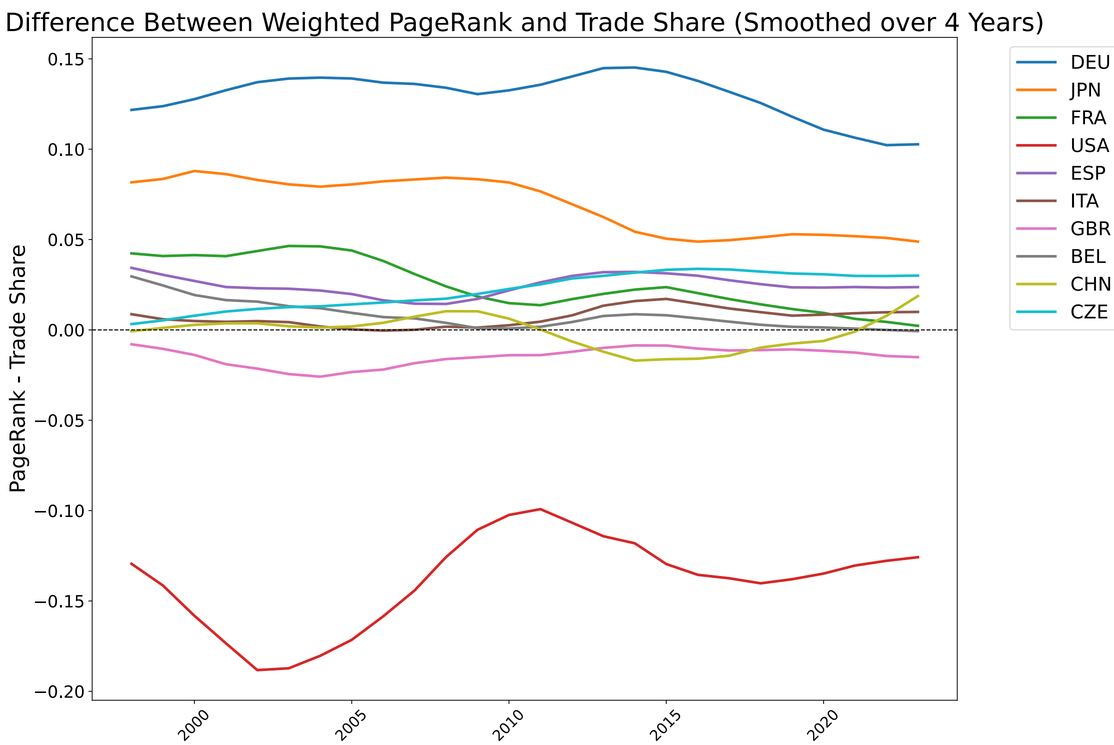

Figure 43: Weighted PageRank - Export Share for top 10 countries of the vehicles trade network. PageRanks are computed using the transpose of the network's weight matrix. $\alpha = 0.99$

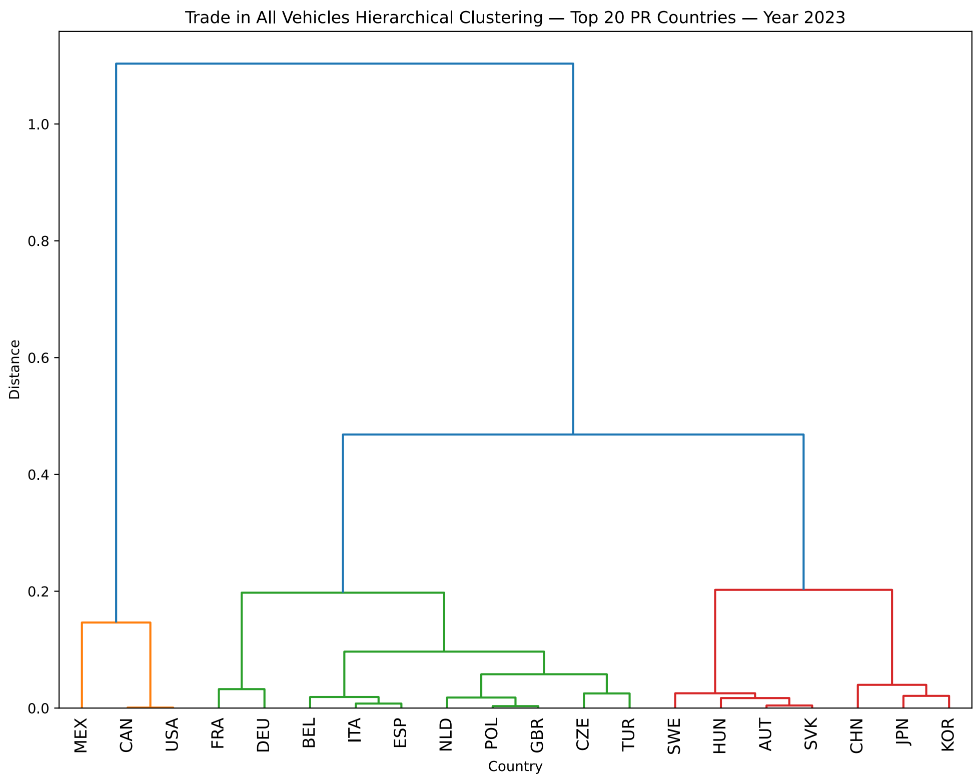

Figure 44: Hierarchical Clustering in 2023 using the eigenvector corresponding to λ_2 for the 20 countries with the highest average weighted PageRanks.