Structure and Properties of Trade Networks [16]

George Hokkanen, Arun Prasad, Ellery Wulczyn

December 9, 2012

Abstract

We analyze the structure of *trade networks*, in which nodes represent agents and edges represent the resource exchanges between them. We study these networks by way of two real-world networks: the *international trade network* (ITN), which contains pairwise trade data between countries over a 50-year period, and the *Commodity Flow Survey*, which contains pairwise trade data between US states over a 5 year period. We describe our insights into the structure of these networks and the extent to which aspects of these structures can be modeled and predicted. We find that a country's GDP can be well modeled by its trade behavior in the ITN. We also beat a canonical model of trade flow prediction and observe interesting relationships between resource trade and population as well as land area in the US.

1 Introduction

Directed graphs have long been used to represent various kinds of social interaction, such as social sentiment among members of a community [11]. We extend this analysis by examining a *trade network*, a more complex case of social interaction where members are self-interested agents who interact with each other in order to increase their total value.

We examine these networks by way of the *international trade network* (ITN), a collection of pairwise trade statistics for all countries from 1948 to 2000. Where applicable, we compare the ITN to the *Commodity Flow Survey* (which we call the STN), a more limited dataset of pairwise trade between US states in 2002 and 2007. For each network, we have total trade values and select commodity trade subgraphs.

We chose these networks because they both exhibit trade patterns in an essentially self-contained way. This is in contrast to datasets like the Prosper moneylending network, in which value (i.e. money) constantly enters the network and relationships change dramatically over time.

Unfortunately, the dynamics of real-world trade networks are complex and not entirely observable. Trade patterns cam emerge from a variety of factors unrelated to economic self-interest, including alliances with other agents, political and social disruption within some aligned group, or wholly external factors like disease and natural disaster. Therefore, we focus on the following two questions:

- *To what extent can trade networks be modeled?* We focus on modeling non-trivial properties of the ITN and try to extend this analysis to the STN.

- *What are the characteristics of trade networks?* This is closely related to the question above, but this question also considers more global features in the network.

Both of these are part of a larger question related to real-world trade networks: how does an agent's trading profile reflect their underlying characteristics? By understanding this relationship, we can describe the characteristics of successful traders and learn how the innate characteristics of an agent determine its actions.

1.1 Data Collection

The ITN data comes from two sources. Total pairwise trade from 1948 to 2000 comes from [7], which is the standard dataset for investigating the ITN, and commodity trade data was scraped from [15], which has comprehensive data for the past 30 years or so despite some significant omissions. All STN data comes from [1].

2 Prior work

Analysis of the ITN has been done in many forms. Some results of the past include demonstrating that GDP and total export are associated by a power law [8] and that the distribution of total trade per country follows a power law. Other experiments have included use of weighted HITS [16] to examine the most influential trading hubs, and graph generation models which rewire a theoretical ITN [3] to test whether or not an emergent property of the ITN is happenstance or by a particular construction. Many of these sources use the same dataset [7] and similar network measurements (e.g. average clustering, average nearest neighbor degree, HITS scores) to the ones we have chosen to use. Despite this, there are interesting properties of the networks left to be described.

The STN data is largely unexplored in a formal context, probably because the dataset is so small. We use the STN data to complement our analysis of the ITN by providing a set of smaller, more comprehensive graphs over many commodities and short time. This allows us to analyze properties of sectoral composition which is difficult with the Comtrade data.

3 Summary Statistics

Here we present an overview of the structure of the ITN.

3.1 Assortativity and Clustering

It has been shown [6] that countries with high degree in the ITN tended to share edges mostly with countries of low degree, demonstrating that nodes were anticorrelated by degree. However, using the weighted construction [3] it was shown that nodes of large weight tend to cluster, demonstrating that the hub nodes of the graph are some of the most tightly connected.

Following these results, we computed the assortativity of the ITN and found it to be between -0.19 and -0.33 based on degree, agreeing with the literature (see figure 1). Something that the literature had not before noted was that the trend of the assortativity of the ITN over the past sixty years has been downward. The average clustering of the ITN is found to be very high, roughly between 0.7 and 0.8, which confirms that the trade network is very dense. While the assortativity decreases over time, the graph's density grows from 0.38 to 0.55. Compared to random graph generation models $G(n, m)$ and preserved degree sequences, we find that this assortativity is not by chance (fixing the number of nodes and edges, and even possibly the degree sequences) since the random models consistently have a mean very near 0. In comparison, commodity networks in the US State trade data have assortativities between -.32 and .23, with a mean of -.05. The aggregate STN has an assortativity of -.11, and it seems reasonable that this is in part due to there being less economic disparity between states in the US than countries around the world.

3.2 Degree Distribution

Over the last three decades of our data (1970-2000) we notice a trend in the degree distribution as the mean shifts rightward and a spike occurs at the most well-connected nodes (figure 2). This coordinates naturally with the increasing density, which would imply that if half of all edges were extant then a great number of nodes would have very large connectivity.

3.3 Reciprocity

We were curious to what extent import and export were correlated between pairs of countries

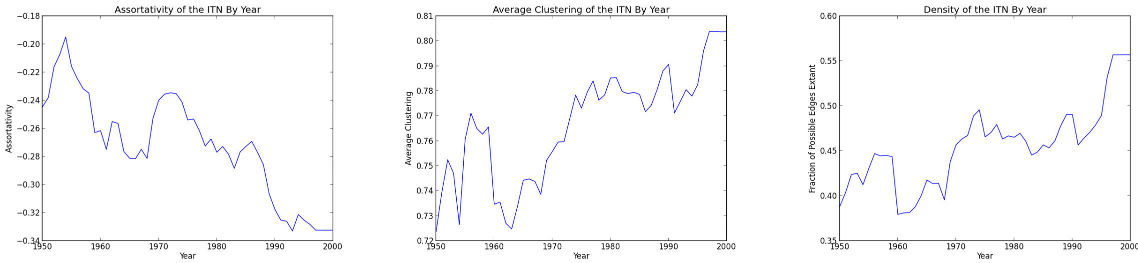

(a) Assortativity (b) Clustering (c) Density

Figure 1: Three Summary Statistics over 1950-2000

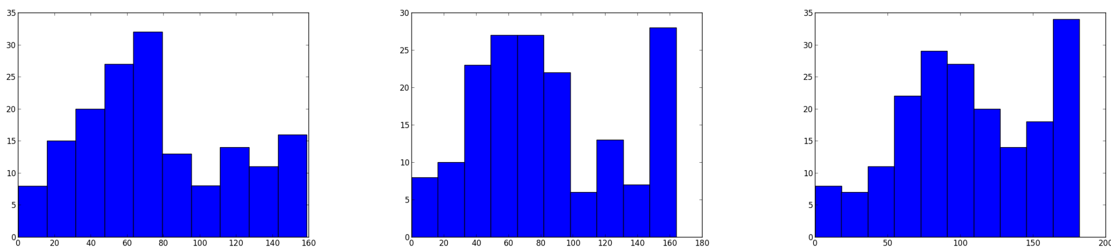

(a) 1980 (b) 1990 (c) 2000

Figure 2: Degree Histograms Per Year

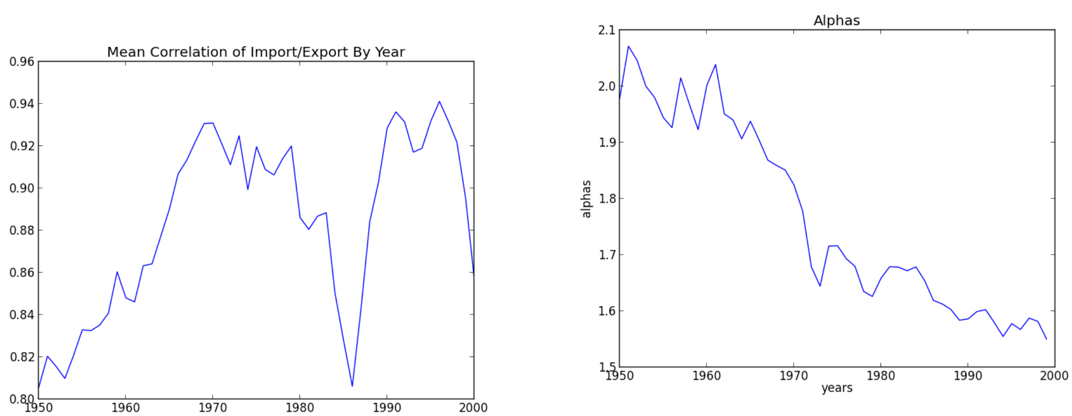

(a) Reciprocity Correlations

(b) Power Law Approximated Exponent over Time that had a directed edge in each direction. We found that reciprocation (the correlation between import and export) is indeed quite high, averaging between .8 to .95 over the time period.

3.4 Clustering

We ran k-means clustering on the set of nodal features. When the data is reduced to its first two PCA scores and visualized in two dimesions, we find that there are three distinct clusters of countries. The first contains highly economically developed countries such as the USA, UK, Japan, Germany, Canada, China, France, Italy, and the Netherlands. The second conatains developed countries such as Portugal, Taiwan, Mexico, Russia, Denmark, etc, and the third contains developing countries such as the Congo, Nepal, Jamaica, Zambia, Yemen, etc.

4 US State Trade Network

4.1 Hypothesis

We hypothesize that some resources are more likely to be produced by states with lower population density (in particular, rural states), while others are more likely to be produced by more urbanized states. For example, we think that a state with lower density would be more likely to have a sectoral composition biased towards producing agricultural products, while one with higher density may instead be producing manufactured goods. Similarly, states with higher overall population would be more likely to be more influential overall in the STN than those with lower overall population, regardless of density. We begin this inquiry by computing correlations between sectoral composition as well as pagerank for each commodity across states, correlating these with population and density.

4.2 Analysis

We perform an analysis on the STN commodity networks for 43 different commodities and observe the correlation of several quantities with respect to total population, population density, and land area. The first of these is the weighted PageRank of each state. The second is $e_s^r = \frac{E_s^r}{E_s}$ where E_s^r is the export of resource r by state s, and $E_s = \sum_r E_s^r$. Similarly we compute this for imports I_s^r. Letting the pagerank of state s in the graph for resource r be pr_s^r and the population of state s be p_s, the correlations reported are $Corr([e_s^r]_S, [p_s]_S)$, $Corr([i_s^r]_S, [p_s]_S)$, and $Corr([pr_s^r]_S, [p_s]_S)$ over all states S for both the graph and its reverse. We similarly compute for the population density d_s and area a_s in place of p_s.

4.3 Findings

Strongly correlated with total population are the total export, total import, summed pageranks in both graphs (these range from .78 to .96), PageRank in the export and import graphs of machinery and furniture, as well as the ratio of machinery to total imports and exports.

Granted, many export graph PageRanks are highly correlated with total population (since more people means more production), but the fact that the highest ones line up very neatly with the sectoral composition (which is a ratio, removing the information with respect to total volume) indicates that the production of machinery has something to do with a state having more people. The most immediate conclusion is that a more populous state will have a more robust economy, which allows for the production of more complicated products.

PageRanks least correlated with total population include the exports of coal, gasoline, and other fuel oils ($< .2$). Compositional indicators most negatively correlated include coal and fuel oil imports ($< -.15$). Also included are cereal grain imports and exports and wood product exports ($< -.2$).

These results demonstrate that how productive a state is in the fuel business tells very little about their population, which might be interpreted to mean that all states, regardless of economy or population, are equally likely to be harvesting these resources. Similarly, the ce-

	sector import	sector export	import pagerank	export pagerank
population	Machinery (.86) Transportation Equipment (.67)	Machinery (.45)	Furniture (.87) Machinery (.86)	Furniture (.97) Machinery (.94)
density	Paper (.63) Transportation Equipment (.46)	Furniture (.59) Transportation Equipment (.48)		
area	Fuel Oils (.56) Gasoline (.55)	Fuel Oils (.47)		Milled Grains (.56)

real grain and wood product compositions of a state's exports are negatively correlated, which is reasonable given that these would more likely be produced by rural states.

Relatedly, we find that fuel oil and gasoline export compositions both have weak correlations ($.5 - .6$) with a state's area. We can take this to mean that prior to observing a state's geography, increasing its size will necessarily increase its oil reverves. The commodity graph in which PageRank has the highest correlation with the state's area is the milled grains network, which is reasonable since such an agricultural product requires a lot of land to be a big producer (and therefore have high PageRank).

5 Predicting GDP based on Structure in the ITN

Studies on the United States, Canada and China have shown a causal relationship between the growth rates of foreign trade and GDP [12] [17]. Since these studies are restricted to only a few countries with large economies, we decided to examine if the relationship between GDP and foreign trade holds more generally across all countries. Instead of attempting to establish a causal realtionship, however, we simply predict a country's GDP based on features computed from the ITN using regression analysis. We chose features that measure either a country's level of trade activity, the level of activity within its trade neighborhood, or the gloabl level of importance as an importer or exporter. These three types of features capture a sense of how a country trades and we will show how this behavior is related to GDP.

We computed the following network features for each country:

1. total degree
2. weighted sum of outgoing edges
3. weighted sum of incoming edges
4. number of triangles
5. local weighted clustering coefficient
6. PageRank
7. weighted PageRank
8. HITS authority score
9. HITS hub score

These are standard network features, but it is worth describing their meaning in the context of the ITN. The total degree is the total number of trade partners a country has. It is the sum of the total number of import and export partners. The weighted sum of outgoing edges is the total export volume, while the weighted sum of incoming edges is the total import volume. The number of triangles measures how often a country's trade partners trade with each other. The local weighted clustering coefficient measures the strength of trade flow within a country's triangles. Both the number of triangles and the local weighted clustering coefficient are measures of how tightly connected the ITN is in the neighborhood of a country. Weighted PageRank in the context of the ITN measures a country's relative importance as an importer. Countries receive high PageRank scores if they import from many countries that are important importers. A measure of a country's relative importance as an exporter is given by its PageRank score on the reversed ITN. The HITS hub score measures how much a country is exporting to important importers and is an alternate measure of a country's importance as an exporter. The HITS authority score measures how much a country is importing from important exporters and is a measure of a country's importance as an importer.

Examining pairwise correlations between GDP and the above network features shows the degree to which there is a linear relationship between the features. We measured a country's trade activity using its total degree, the sum of its imports and the sum of its exports. The degree, which captures only information on the amount of trade relationships and not the amount of trade volume, is positively correlated with GDP (0.34), but only weakly so compared to total exports (0.84) and total imports (0.89). When we look at the features that measure trade ativity in a country's neighborhood, we see a similar result. The number of triangles, which also does not account for trade volume, is weakly positively correlated with GDP (0.35), while the local weighted clustering coefficient is highly correlated with GDP (0.77). Global measures of importance as an exporter or importer show high correlations with GDP. Given that, for all features that a have a weighted and an unweighted version, the weighted version shows higher correlations with GDP, it is surprising that the HITS authority score, which does not use edge weights, has the highest correlation with GDP (0.90) of all.

Looking at pairwise plots of GDP with other features suggests that the relationship between GDP and total import, total export and PageRank is not linear. Taking log transformations of all the variables shows that correlations with GDP go up signifiantly for the unweighted features degree, triangles and PageRank, although they remain lower than their weighted analogs. The correlations between GDP and total import, total export and HITS hub scores also increase under the log transform.

For both the regular and transformed data sets, we performed the following regression procedure. We first center and scale the data (i.e. every feature has mean zero and variance 1) such that errors in the regressions on the two different data sets are comparable. Then we run an all-subsets selection algorithm to chose the subset of features with the lowest AIC score. We compare the two models using cross-validation test error to determine the better model.

	Coefficients
(Intercept)	0.00
	(0.01)
w-pr	1.27^{***}
	(0.35)
T import	-2.62^{***}
	(0.48)
T export	0.90^{***}
	(0.10)
pop	0.35^{***}
	(0.01)
hubs	-0.11^{***}
	(0.03)
authorities	1.44^{***}
	(0.11)
R^2	0.97
Adj. R^2	0.97
Num. obs.	184

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 2: Untransformed features chosen from all-subsets selection

For the untransformed data, all-subsets regression chose weighted PageRank, total import, total export, population, hubs score, and authorities score as features. Regressing on this subset gives an $R^2 = .97$, and the coefficients for each feature are highly significant. For all features, except for total import and hubs score (importance as an importer) the coefficients are postive. This indicates that ceteris paribus, having higher imports or being a country that imports more from important exporters is related to lower GDP according to this model. The inverse holds for having higher exports and exporting more to important importers. We did a 10-fold cross-validation procedure to compute

transfrom	degree	export	import	triangles	clustering	pr	wpr	authortity	hub
log	0.67	0.90	0.90	0.65	0.74	0.81	0.82	0.76	0.73
none	0.34	0.84	0.90	0.35	0.77	0.34	0.88	0.90	0.52

Table 1: Correlations with GDP

	Coefficients
(Intercept)	0.00
	(0.01)
weighted PageRank	0.17
	(0.11)
total import	0.50^{***}
	(0.10)
degree	0.57^{***}
	(0.15)
total export	0.24^{***}
	(0.07)
pop	0.47^{***}
	(0.02)
PageRank	0.30^{***}
	(0.06)
authorities	-0.22^{**}
	(0.09)
triangles	-1.09^{***}
	(0.15)
R^2	0.96
Num. obs.	184

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 3: Log-transformed features chosen from all-subsets selection

the test MSE, which is 0.47.

The U.S. is a major outlier and the fold that was not trained on the U.S. data point gets a large error. If we remove the U.S. from the data set and rerun the all-subsets regression we get a very different subset of features and a much lower crossvalidated MSE of 0.23. This shows how unstable the model is due to this extreme outlier point. Using the log-transformation of the data leads to a much more stable model.

The number of triangles, PageRank and degree enter the optimal feature set for the log-transformed data. It is striking that these are all unweighted features and that the correlations with these unweighted features reach into the 0.80's. The R^2 value is a bit lower at 0.96 but the cross-validated MSE is an order of magnitude lower at 0.042. We can compare the MSE across transformations of the data set since the data was centered and scaled. When removing the U.S., the optimal subset of features and cross-validated MSE are unchanged.

In conclusion GDP can be closely modeled by using network features from the ITN. Features from each of the three classes described above significantly contribute to the models. Measures of a country's trade activity, the level of trade activity in its neighborhood and the global importance of a country as an importer or exporter are all useful in modeling GDP. Without transfromations to the data, features that use edge weights are the most powerful, but making a log transformation makes the unweighted features important as well. Finally a log transformation leads to model that is less susceptible to outliers and has better test error.

5.1 Predicting Trade Flows

A basic model used for predicing bilateral trade flows between countries in the extant literature [2] is the gravity model:

$$w_{ij} = \left(\frac{\mathrm{GDP}_i^a \mathrm{GDP}_j^b}{dist_{ij}^c} \right) \phi$$

According to this model large economies that are close together geographically will trade more. Another way to write down the model is

$$\log(w_{ij}) = a\mathrm{GDP}_i + b\mathrm{GDP}_j + c\mathrm{dist}_{ij} + e$$

This says that there is a linear relationship between the logarithm of the trade between

	Coefficients
(Intercept)	58.53^{***}
	(0.30)
PageRank exporter	4.25^{***}
	(0.05)
authorities exporter	1.24^{***}
	(0.02)
PageRank importer	6.14^{***}
	(0.04)
R^2	0.60
Num. obs.	33856

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 4: Best Three Feature Model

two countries and the logarithm of the GDPs and the distance between the two countries. A limitation of this model is that it does not differentiate between import and export since $w_{i,j} = w_{j,i}$. Running a regression on the ITN data shows that this simple model gets an R^2 value of 0.55. The coefficients on GDP are postive and the coefficient on distance is negative confirming the intuition that large economies that are close geographically will trade more.

Our second modeling goal was to extend the gravity model using features from the ITN. To build these models we assume that we have an unweighted version of the ITN at our disposal. This means that we know which countries trade with one another but we do not know how much. If an edge had weight zero in the weighted graph, it is removed from the unweighted graph, but we cannot tell the difference between a trade flow of \$1 and a trade flow of \$100,000,000. The features at our disposal are the distance between countries and the unweighted country features described above.

This includes the GDP, degree, number of triangles, HITS hubs and authorities, and PageRank. The best model using only three features like the gravity model includes PageRank the exporting and importing country and the authority score of the importing country. This model achieves an R^2 of 0.60. It has the advantage, that it can differentiate between an import and an export and that it uses knowledge of whether countries trade at all. If we undo the log transformation, then the model says:

$$w_{ij} = 58.53 pr_i^{4.25} pr_j^{6.14} \mathrm{auth}_i^{4.25}$$

So the higher the PageRank of the two countries and the higher the authority score of the exporting country, the greater the volume of export. It is interesting that the best features turn out to be measures of global importance as an importer or exporter, like GDP in the gravity model, but that that the distance feature, which is the only data specific to an edge, is omitted in favor of another feature that measures global importance. It would seem the Gravity model has become outdated in a globalized economy.

If we do not restrict ourselves to a subset of size three for sake of comaparison to the gravity model, and choose the subset of models with the lowest AIC score, we get a model with R^2 of 0.65.

6 Conclusions and Further Work

We have provided models to outperform canonical trade modeling, found that a country's GDP is strongly related to its trade profile, and observed correlations between population and area in commodity trade. For future work, we would ideally compile a composite dataset witnessing the best properties of both the ITN and STN. Such a dataset would have a fine commodity granularity and long temporal history, and may even demonstrate both micro- and macro-level trade structure (intra- and international). Given such data, one could model the evolution of a trade network and model trade flows by comparing the resource production profiles of countries. This could lead to an empirical theory of how countries acquire resources that they do not produce.

References

[1] U.S. Census Bureau. Commodity flow survey, 2002, 2007.

[2] Giorgio Fagiolo. The international-trade network: gravity equations and topological properties. *Journal of Economic Interaction and Coordination*, 2010.

[3] Giorgio Fagiolo, Javier Reyes, and Stefano Schiavo. On the topological properties of the world trade web: A weighted network analysis. *Physica A: Statistical Mechanics and its Applications*, 387(15):3868–3873, 2008.

[4] Organisation for Economic Co-operation and Development. Main economic indicators. http://stats.oecd.org/.

[5] International Monetary Fund. Global data. http://www.imf.org/external/data.htm.

[6] Diego Garlaschelli and Maria I. Loffredo. Structure and evolution of the world trade network. *Physica A: Statistical Mechanics and its Applications*, 355(1):138–144, 2005.

[7] Kristian Skrede Gleditsch. Expanded trade and gdp data. http://privatewww.essex.ac.uk/ ksg/exptradegdp.html.

[8] J Saramaki K Kaski K Bhattacharya, G Mukherjee and S S Manna. The international trade network: weighted network analysis and modelling. *Journal of Statistical Mechanics*, 2008.

[9] Tim Kastelle. Generative mechanisms of the world trade web. In *DRUID Summer Conference*, 2009.

[10] Taedong Lee and Byoung-Inn Bai. Do birds of a feather trade together? homophily in preferential trade agreements network. In *International Political Economy Society Conference*, 2009.

[11] Jure Leskovec, Daniel P. Huttenlocher, and Jon M. Kleinberg. Signed networks in social media. In Mynatt et al. [13], pages 1361–1370.

[12] Xiang Li, Yu Ying Jin, and Guanrong Chen. Complexity and synchronization of the world trade web. *Physica A: Statistical Mechanics and its Applications*, 328(1–2):287–296, 2003.

[13] Elizabeth D. Mynatt, Don Schoner, Geraldine Fitzpatrick, Scott E. Hudson, W. Keith Edwards, and Tom Rodden, editors. *Proceedings of the 28th International Conference on Human Factors in Computing Systems, CHI 2010, Atlanta, Georgia, USA, April 10-15, 2010*. ACM, 2010.

[14] Frank Schweitzer, Giorgio Fagiolo, Didier Sornette, Fernando Vega-Redondo, and Douglas R. White. Economic networks: What do we know and what do we need to know? *Advances in Complex Systems*, 12(04n05):407–422, 2009.

[15] United Nations Statistics Division. United Nations Commodity Trade Statistics Database (UN comtrade). http://comtrade.un.org/.

[16] Gang Liu Wei Wei. Bringing order to the world trade network. *International Proceedings of Economics Development and Research*, 28:88–92, 2012.

[17] Tao Xiangnan Zestos, George K. Trade and gdp growth: Causal relations in the united states and canada. *Southern Economic Journal*, 2002.