An Examination of Alternative Multiplex Modularity Maximization Models

Jason Cheal

Advised by Scott Pauls

May 22, 2019

Presented to the faculty of the Department of Mathematics

Dartmouth College

In partial fulfillment of the

requirement for the degree of the

Bachelor of Arts

Abstract

Multiplex networks are versatile objects capable of representing many complex systems. While offering increased flexibility and representative power, extending single layer analytical techniques to the multiplex setting presents many questions. We study the extension of modularity maximization, a community detection strategy that is well understood for single layer networks, to the multiplex case. We consider two alternative models (diagonal and non-diagonal) and examine the behavior of each on different families of networks. We first apply each model to a series of toy networks with well understood characteristics, and then apply each to World Trade Web (WTW) networks. Findings from the toy case analysis indicate that both models consistently identify robust embedded global communities and increasingly favor layer communities as the communities become more different across layers. The WTW analysis mirrors these results, indicating that these findings bear out in the analysis of real world systems. Additionally, the WTW analysis indicates that the two models emphasize different multiplex network characteristics, with the diagonal model showing particular sensitivity to the alignment of heavily weighted layer communities, and the non-diagonal model more sensitive to the similarity of node copy neighborhoods across layers. We justify and concretely state each of these conjectures, thus helping to inform the contexts in which each model may be most appropriate in future research.

I. Background

Networks are a nearly ubiquitous object within the field of applied mathematics. This is primarily due to their ability to capture the characteristics of highly complex systems, such as social networks, the World Wide Web, or transportation infrastructure (Banerjee et. al., 2013; Fagiolo et. al., 2010; Jenelius et. al., 2006). Multiplex networks are a class of networks composed of a single set of nodes, k set of edges, and k functions relating each edge in a given edge set to a pair of vertices, i.e. networks with multiple layers, each containing the same vertex set but different edges (Menichetti et. al., 2014). Multiplex networks generalize single layer networks and allow for the representation of more granular characteristics of a system. For example, when examining social networks, multiplex networks allow for a more nuanced representation, with different types of social relationship on each layer

(represent family connections on one layer, professional relationships on another, etc.). In this way, multiplex networks open countless new avenues for research. But, such explorations require tools of analysis, many of which can be extended from those applied to single layer networks. Modularity maximization is an example of one such tool.

Conceptually, modularity maximization is an analytical strategy that aims to identify highly connected sub-communities within a network system. For general background on modularity maximization, see Newman (2010, 373-382). To identify these communities, the configuration model is used, in which cluster connectivity is compared to the connectivity of an Erdös-Rényi random network of the same degree distribution. Specifically, the configuration model is constructed by cutting each edge in the network in half and randomly rewiring the resulting stubs. For a single layer network *G* with a given degree sequence, the resulting configuration model *C* will display the following property:

$$P[(i,j) \in E(C)] = \frac{d_G(i)d_G(j)}{2e(G)}$$

Using this, we build the adjacency matrix *R* for the random configuration model as follows:

$$R_{i,j} = \frac{d_G(i)d_G(j)}{2e(G)}$$

The modularity matrix *B* is the difference between this random configuration matrix and the adjacency matrix:

$$B_{i,j} = A_{i,j} - R_{i,j}$$

By construction, $B_{i,j}$ represents the connection of node *i* to node *j* in comparison to what might be expected in a randomly generated network of the same degree distribution.

Next, we define an indicator vector s for a candidate partition with communities C_1 and C_2 as follows:

$$s(i) = \begin{cases} 1 \; if \; i \in C_1 \\ -1 \; if \; i \in C_2 \end{cases}$$

Using this vector, we can calculate the difference between intra-community connectivity in the network of interest and the corresponding configuration model for a given partition, which is referred to as the modularity Q:

$$Q = \frac{s^T B s}{2e(G)}$$

Maximizing the modularity identifies the most highly connected sub-communities within the network structure. Additionally, this strategy has a natural termination, as no further subdivision is valuable when the change in modularity resulting from any possible further subdivision is less than or equal to 0, i.e. $\Delta Q \leq 0$.

A relaxed optimization of the preceding equation is performed by redefining s to be a unit vector with entries in $\mathbb{R}$ and calculating $argmax_{|s|=1} s^T B s$. Using Lagrange multipliers, the vector s that maximizes Q is a unit eigenvector associated with λ_n, the largest nonzero eigenvalue of B. For a more detailed description of this relaxed optimization see Newman (2010, 376-377). This unit eigenvector s determines the partition, with all nodes of positive sign placed in C_1 and all those with negative signs placed in C_2. This process can be iterated, creating more and more sub-communities until the change in modularity resulting from any possible further sub-division is negative.

Modularity maximization is quite straight forward conceptually. However, extending the strategy to the multiplex case raises many questions. For example, any reasonable extension must scan both within and between layers to identify highly connected sub-

communities that may span multiple layers. Consider, for example, a social network containing a group of colleagues that are also closely knit socially, and thus highly connected on two different layers. Any reasonable multiplex modularity maximization model should cluster these nodes across both layers, as the group represents a tighter community than any group of people that are either just friends or just colleagues. Thus, convention must be established regarding how inter-layer connections are represented and treated analytically.

One possible approach, which we will refer to as the "diagonal model," treats edges between node copies as the only inter-layer connection. This representation calls for a block diagonal modularity matrix $\hat{B}_{i,j}^{\alpha,\beta}$ in which the diagonal blocks contain layer modularity matrices equivalent to the single layer case described above. The off-diagonal blocks are multiples of the identity matrix, with a weight parameter *w* that varies the weight of edges between node copies:

$$\hat{B}_{i,j}^{\alpha,\beta} = \begin{cases} B_{i,j}^{\alpha} \; if \; \alpha = \beta \\ w \; if \; \alpha \neq \beta \; and \; i = j \\ 0 \; otherwise \end{cases}$$

Applying modularity maximization yields a partition of the multiplex network into highly connected communities that may span multiple layers.

While the diagonal model is straight forward, it only allows node copies to have an unusually high inter-layer connectivity, as seen in the off-diagonal blocks of the modularity matrix being multiples of the identity. One might expect this to hinder the ability of modularity maximization to identify sub-communities that span layers. An alternative method, which we will refer to as the "non-diagonal model," places inter-layer edges between v_i^{α} and v_i^{β} if either $v_i^{\alpha} \sim v_j^{\alpha}$ or $v_i^{\beta} \sim v_j^{\beta}$, i.e. node *i* on layer α is adjacent to node *j* on layer β if the two nodes are adjacent on either layer (throughout paper, v_i^{α} will refer to node *i* on layer α).

In execution, these inter-layer edges are formulated using a mixing matrix *M.*

$$M_{i,j}^{\alpha,\beta} = \begin{cases} mc \; if \; i = j \\ 0 \; otherwise \end{cases}$$

In the above definition, m represents an amplifying constant which is equal to 1 throughout all proceeding analysis. The constant c represents a distribution constant that conceptually represents the proportion of the information at n_i^α that is transferred to its copy n_i^β. Throughout all proceeding analysis c is equal to $1/k$, i.e. information spreads equally between all layers. This is the equidistribution model, the simplest example of non-diagonal multiplex models (DeFord and Pauls, 2017).

This mixing matrix is multiplied on the left to the block diagonal adjacency matrix $\hat{A}$ and the block diagonal configuration matrix $\hat{R}$, defined as follows:

$$\hat{A}_{i,j}^{\alpha,\beta} = \begin{cases} A_{i,j} \; if \; \alpha = \beta \\ 0 \; otherwise \end{cases}$$
$$\hat{R}_{i,j}^{\alpha,\beta} = \begin{cases} R_{i,j} \; if \; \alpha = \beta \\ 0 \; otherwise \end{cases}$$

The resulting mixed configuration matrix is subtracted from the mixed adjacency matrix to obtain the modularity matrix $\hat{B}$:

$$\hat{B} = M\hat{A} - M\hat{R}$$

Once again, the iterated relaxed optimization is then applied to this modularity matrix to obtain the optimal partition. For further background, justification, and explanation of these two multiplex formulations, see DeFord (2018, 61-64).

The preceding discussion has focused most explicitly on undirected, unweighted networks; however, the conceptual framework equally applies to directed and weighted cases. The undirected case is simply a subset of the more general directed case in which two directed

edges exist between any connected nodes. Loosening the model to allow one-way connections results in a non-symmetric modularity matrix that must be symmetrized prior to eigenvector decomposition, and node in-degrees and out-degrees must be averaged when calculating modularity, but the analysis is otherwise analogous. Conceptually, there is little to indicate that this adaption to more complex directed systems will result in meaningful analytical differences, though this will be tested in practice in the subsequent sections.

For the weighted case, edges, and thus super-adjacency matrix entries, are no longer binary but instead convey information regarding edge "importance" or "strength." In the configuration model process, edge *e* of weight *d(e)* is split into 2 stubs, each of weight $\frac{d(e)}{2}$. If a stub of edge *e* is rewired with a stub of edge *f,* the resulting edge is assigned weight $\frac{d(e)}{2} + \frac{d(f)}{2}$. While fundamentally congruent to the unweighted case, one might expect the introduction of weights to result in greater variability in results, as edges of unusually large weights can have a large impact on clustering by representing an outlier value in the modularity matrix, a hypothesis that will be tested in proceeding sections.

II. Introduction

In this paper, we examine the characteristics of diagonal and non-diagonal modularity maximization. We seek to build out a set of conjectures regarding the performance of each model on different types of multiplex networks.

In Section 3, we begin developing these conjectures by applying each model to a series of toy networks with well understood embedded community structure. By examining the results of each model when applied to multiplex networks with aligned layer communities, we develop the following conjecture:

Conjecture 1: When applied to multiplex networks with well-defined, robust global communities, both diagonal and non-diagonal modularity maximization identify the global communities and partition nodes accordingly.

Next, each model is applied to two-layer multiplex networks in which the layer communities are offset to varying degrees. The results of these tests provide the basis for our second conjecture:

Conjecture 2: As layer communities are increasingly offset, both diagonal and non-diagonal modularity maximization increasingly favor layer communities over global structure, separating node copies that belong to highly different communities on each layer.

Lastly, we find divergent tendencies between the two models when applied to networks with highly offset layers. In such cases, the diagonal model continues to emphasize layer communities, separating the multitude of node copies that belong to highly different communities on each layer. The non-diagonal model, however, creates more nuanced partitions, identifying pockets of similarity within node copy communities that are highly different across layers and emphasizing global structure to a greater degree than the diagonal model. This is seen in each models' partition classifications at large offset values, with the diagonal model continuing to form communities that emphasize layer communities (and separate many node copies) while the non-diagonal model forms more nuanced, though "illogical," partitions that group together more of the node copies despite their highly different layer communities.

In Section 4, we apply each model to World Trade Web (WTW) data and interpret the results in the context of Conjectures 1 and 2 by building a set of five hypotheses. Specifically, we apply each model to pairs of commodity layers within the WTW. By classifying each commodity pair as *Supply Chain, Similar Consumables,* or *Seemingly Unrelated,* we are able to interpret each model's performance in the "aligned" vs. "offset" framework developed in the toy case analysis. Our findings support Conjecture 1 and 2, as both models favor global structure in cases of similar layer community structure (eg. when

applied to a *Supply Chain* pair) and favor layer community structure in cases of highly offset layer community structure (eg. when applied to a *Seemingly Unrelated* pair). Additionally, the results provide the basis for a third and final conjecture:

> *Conjecture 3:* The diagonal and non-diagonal models emphasize different multiplex network characteristics, with the diagonal model more sensitive to the alignment of heavily weighted layer communities, and the non-diagonal model more sensitive to the similarity of node copies' first and second order neighborhoods across layers.

This tendency is highlighted through a series of tests using two alternative measures of "different-ness," to which the two models react divergently.

Lastly, we conclude in Section 5 by summarizing our findings and the implications for future analysis and research.

III. Toy Case Analysis

To understand the differences between the diagonal and non-diagonal multiplex models, we first examine their performance across a family of networks with explicit planted communities. To do this, we use the Stochastic Block Model, embedding two communities on each layer, with intracommunity connection probability of p and intercommunity connection probability of q (with $p \geq q$). For other examples of this type of generative model, see Stanley et. al. (2016) and Bazzi et. al. (2016). In the first series of tests, these communities are aligned across k layers, ensuring robust global communities. Next, the communities are offset across two layers, such that portions of each community overlap across layers, while other portions are misaligned, leading to ambiguous global community structure.

During these tests, the majority of parameters, as defined in Section 1, are set at default values, listed below:

$$Node\ set\ size\ (n) = 100$$
$$Probability\ of\ intracommunity\ connection\ (p) = 0.5$$
$$Layer\ community\ size = 50$$
$$Diagonal\ method\ node\ copy\ edge\ weight\ (w) = 1$$
$$Non-diagonal\ method\ distribution\ constant\ (c) = 1/k$$
$$Non-diagonal\ method\ amplifying\ constant\ (m) = 1$$

Within the aligned test, we vary the number of layers (k) and the intercommunity edge probability (q). In the offset test, we set k equal to 2 and vary the offset parameter (ℓ) as well as q. The offset parameter determines the robustness of the network's global communities. The aligned tests imply $\ell = 0$, i.e. layer planted communities perfectly overlap. As ℓ increases, the Layer 2 embedded communities are rotated in relation to the Layer 1 embedded communities, forming an increasingly large ambiguous region of nodes that are embedded in different communities on each layer, as the overlapping region shrinks. This allows for the examination of model performance in situations of increasing global community ambiguity.

To evaluate the efficacy of the partition in each test, we examine the resulting *Mistake Ratio* (r), defined as follows:

$$r = \frac{nodes\ mistakenly\ partitioned}{total\ possible\ mistakes}$$

In the above definition, nodes are considered to be mistakenly partitioned if they are not grouped within their planted community. In the aligned case, every node is within a robust global community, so the total possible mistakes is equal to $\frac{nk}{2}$. In the offset case, only nodes within overlapping community regions belong to a robust global community, so the total possible mistakes is equal to $100 - 2\ell$. While this metric does not provide color on how the models handle the ambiguous misaligned nodes in the offset tests, it assesses each

models' ability to identify robust global communities, and thus allows us to test Conjecture 1.

3.1 Aligned Layer Community Tests

The aligned case is a fundamental test of the ability of both modularity maximization models to identify obvious and robust global communities. For any number of layers, and given a q value that is substantially less than p, any reasonable approach to modularity maximization must correctly identify the two global communities. To explicitly test this, we limit the number of communities to two for each model across all aligned trials.

We first perform the aligned test in the case of an unweighted, undirected multiplex network, with the following parameters:

$$q = \{0, 0.05, \ldots, 0.5\}$$
$$k = \{2, 3, 4, 5, 6\}$$

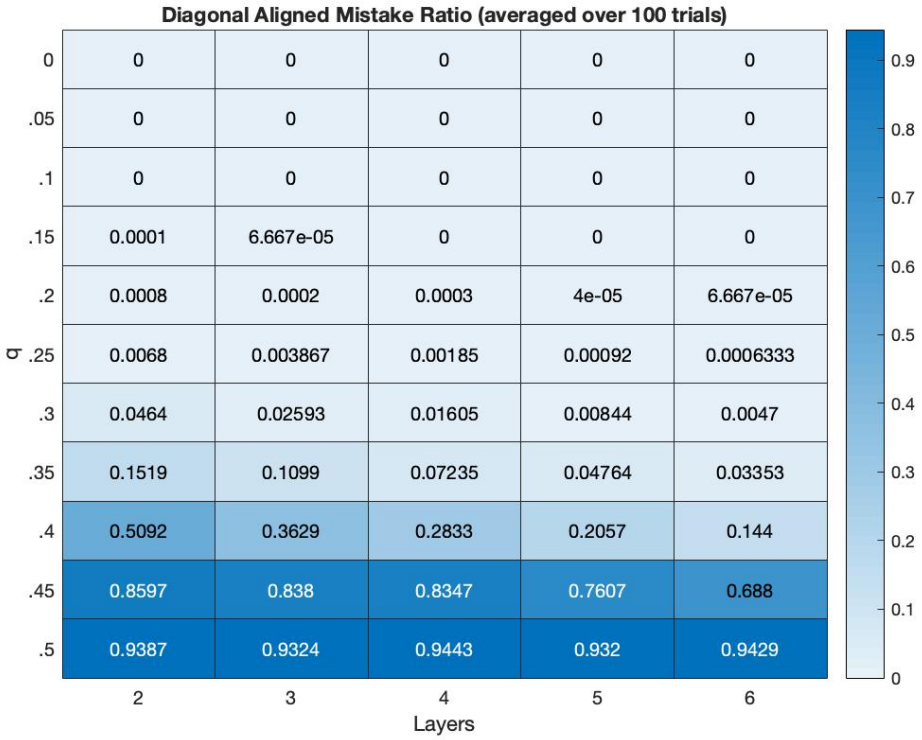

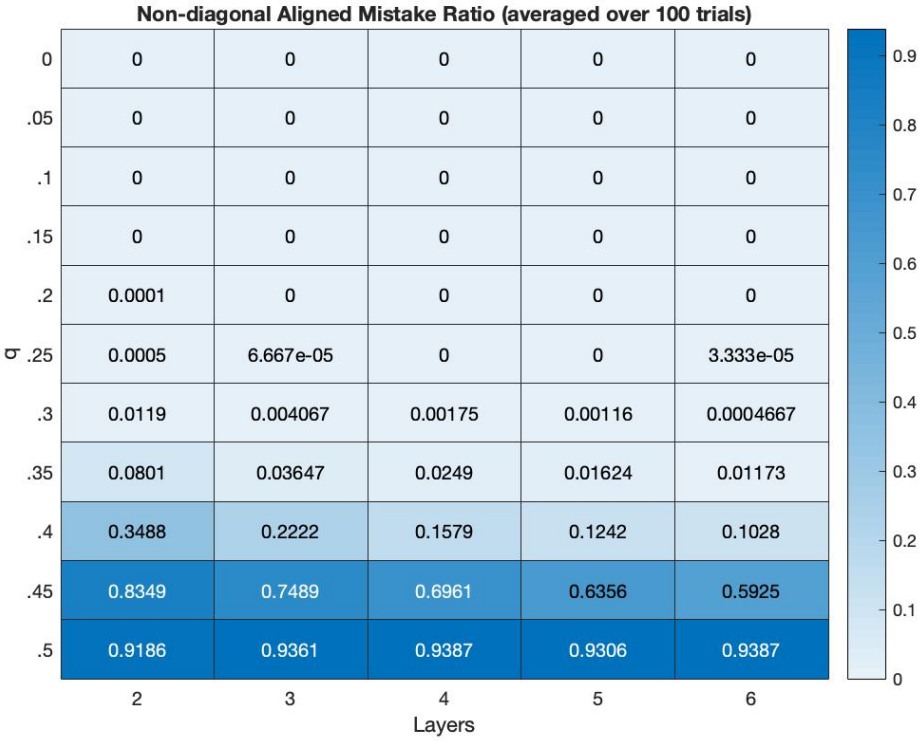

Figure 3.1: Unweighted-undirected aligned test Mistake Ratio results (averaged over 100 trials)

The results of this test, shown in Figure 3.1, indicate the promise of both models. For all $q \leq 0.2$, both the diagonal and non-diagonal analysis perfectly identify the global communities. As q approaches p, both models perform increasingly poorly, with the non-diagonal model somewhat outperforming the diagonal model for cases in which q is only slightly less than p (eg. q=0.4 and q=0.45).

We next perform a more generalized version of the aligned test by making the networks directed and weighted. The weights are assigned randomly from the following distribution D:

$$D = \bigcup_{i=1}^{100} \bigcup_{j=1}^{i} \{101 - i\}$$

In other words, D contains the value 100 repeated once, 99 repeated twice, etc., down to 1 repeated 100 times.

The purpose of this weighted and directed aligned community test is to ensure that each model is still able to correctly identify robust global communities in a more complex system, and to ensure Conjecture 1 holds in these more generalized cases. We leave all other parameters unchanged, and apply both models to networks with varying q and k values, in search of two communities.

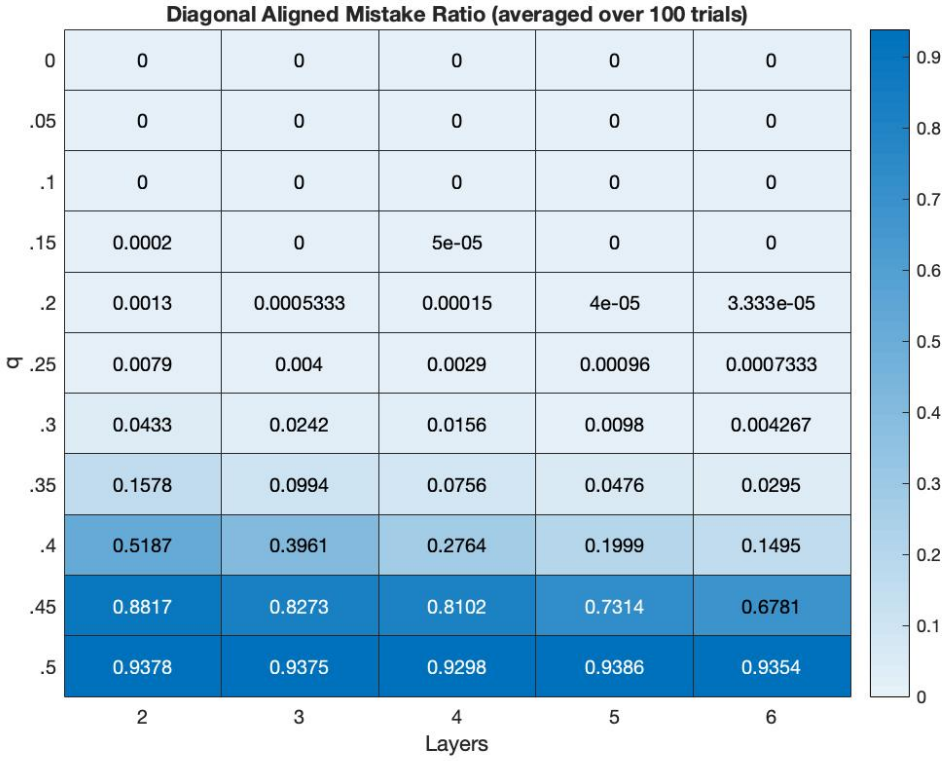

Figure 3.2: Weighted-directed aligned test Mistake Ratio results (averaged over 100 trials)

As seen in Figure 3.2, both models once again appear to pass this baseline test. For all $q \leq 0.3$, both models generate essentially zero mistakes, correctly identifying the global communities. As previously, performance deteriorates as q approaches p. In this test, the non-diagonal model once again appears to somewhat outperform the diagonal model across the board, making slightly fewer mistakes at all levels of $q \geq 0.2$.

Taken together, the results of these aligned layer community tests provide initial support for Conjecture 1, as each model is correctly able to identify robust (aligned) and well-defined (q meaningfully less than p) global communities. As these communities become less well-defined, both models perform increasingly poorly. The fact that the non-diagonal model slightly outperforms the diagonal model is likely related to the handling of inter-layer connections. As the diagonal model only allows for inter-layer connections between node copies, inter-layer communities are muted. However, the non-diagonal model's more nuanced handling of inter-layer edges emphasizes global structure over that of each individual layer. As a result, the non-diagonal model is better able to identify poorly defined communities that are robust across layers, as the inter-layer connections serve to amplify the weak layer community structure.

3.2 Offset Layer Community Tests

The case of embedded communities that are offset across two layers provides a good test of each model's ability to identify global communities in increasingly ambiguous situations. Throughout these tests, we vary the offset value (ℓ), embedding the first layer with communities containing nodes {1,…,50} and {51,…,100}, and embedding communities {ℓ+1,…,ℓ+50} and {1,…,ℓ}U{ℓ+51,…,100} on the second layer. We test values of ℓ varying from 1 (maximal overlap) to 25 (minimal overlap). As in the aligned layers trials, we first apply both models to the case of an unweighted, undirected network, with q={0,0.05,…,0.5}, and search for two communities.

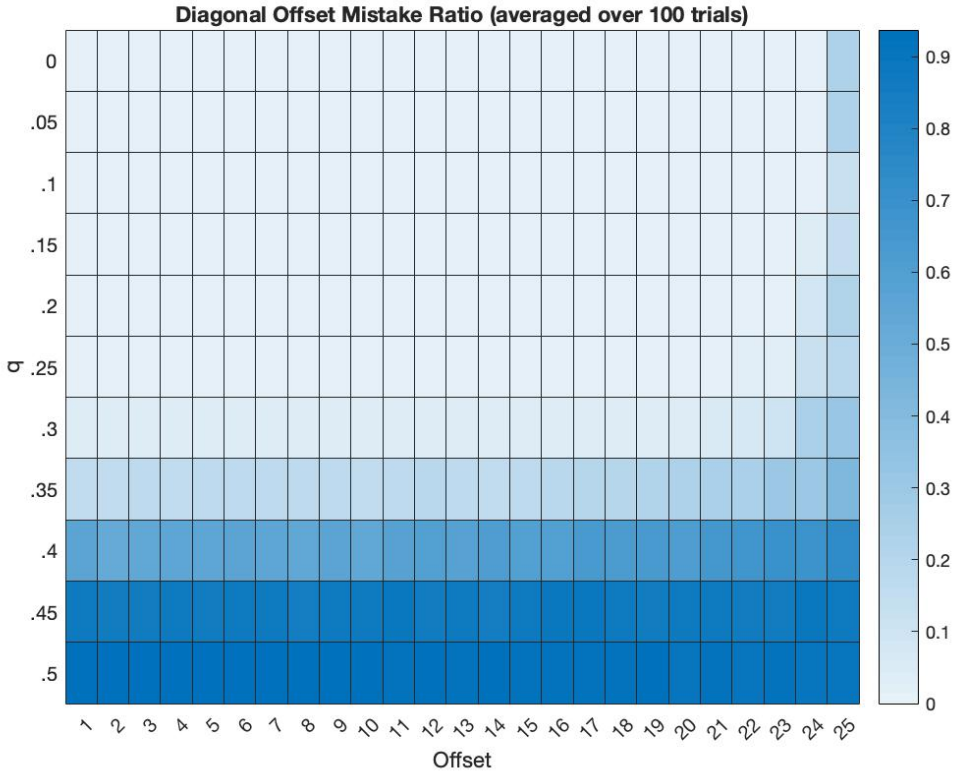

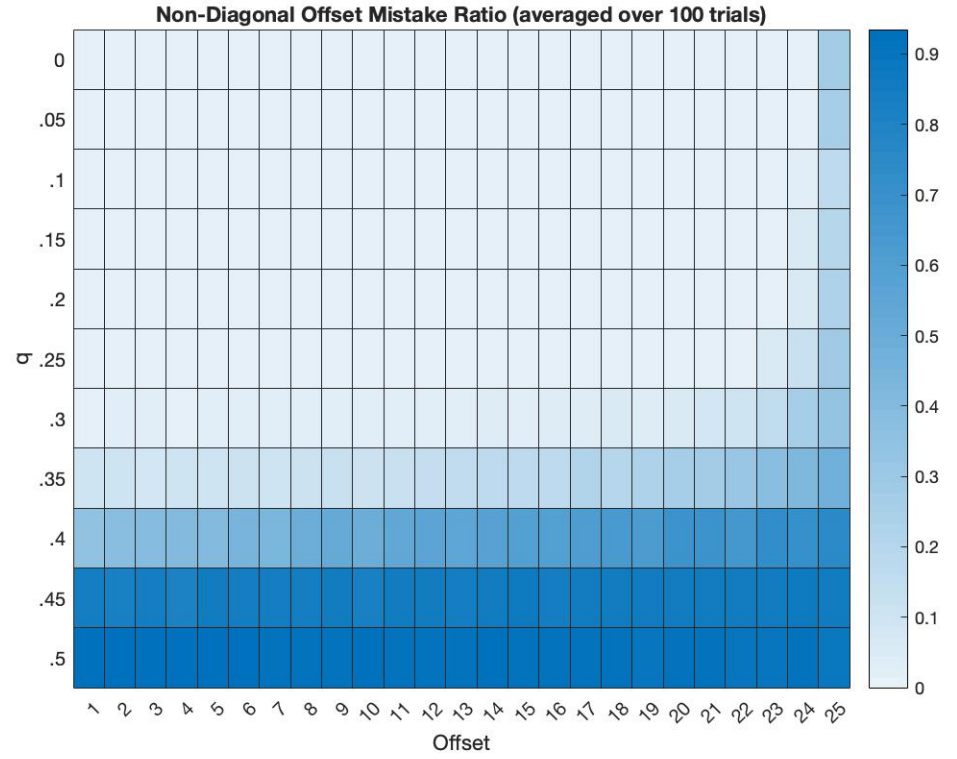

Figure 3.3: Unweighted-undirected offset test Mistake Ratio results (averaged over 100 trials)

Examining these results, shown in Figure 3.3, it is important to note that mistakes are only identified within the overlapping regions. For example, if ℓ =5 only the 90 nodes whose embedded communities overlap across the two layers are examined for mistakes. We will attempt to address this gap in our understanding by systematically examining ambiguous region partitions in Section 3.4 below. However, the *Mistake Ratio* still provides insight and sheds meaningful light on the validity of Conjecture 1.

These results reaffirm that both models perform quite strongly for values of q that are meaningfully less than p. As q approaches p, performance deteriorates for both models, with the non-diagonal model slightly outperforming the diagonal. It is interesting to note that in both cases, the amount of offset has little impact on model performance. This is likely due to the fact that the *Mistake Ratio* only assesses overlapping nodes. This highlights the importance of a more granular examination of how each model treats ambiguous nodes, which is performed in the Section 3.4 below.

Before proceeding, we again generalize the offset community detection test by expanding to the weighted and directed case. As previously, the weights are assigned from the distribution D. In all other respects, this test is identical to the unweighted-undirected offset test as described above.

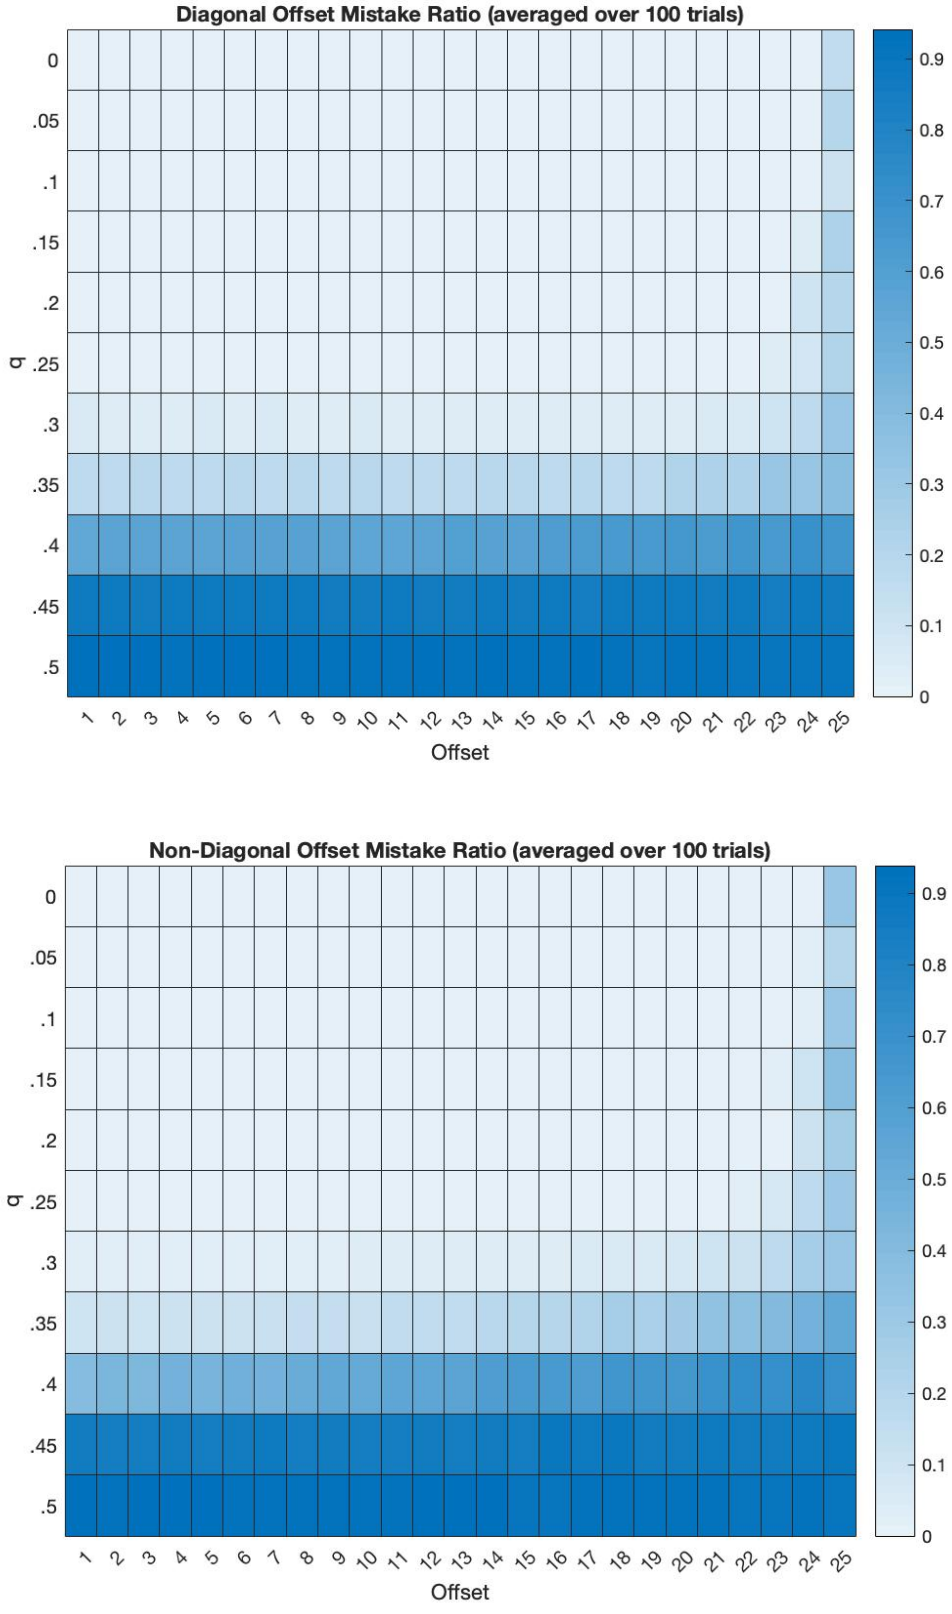

Figure 3.4: Weighted-directed offset test Mistake Ratios (averaged over 100 trials)

The results of this generalized offset communities test, as shown in Figure 3.4 above, are generally similar to the results of the initial offset test. Both models again perform strongly for values of q that are meaningfully less than p. However, both models, and the diagonal model in particular, perform more poorly at the largest offset values. We examine these large offset value results on a case-by-case basis and find that both models exhibit a tendency to group large clusters of nodes from one layer together, spanning across overlapping and ambiguous regions. One potential explanation for this phenomenon is the existence of one or more heavily weighted edges connecting to ambiguous nodes. The results here indicate that layer communities containing heavily weighted edges can dominate the global community partition produced by each model.

We will further explore the effects of heavily weighted edges and layer communities in WTW analysis Sections 4.4 and 4.5. That aside, for now we can conclude that the performance of each model deteriorates rapidly at the extremes of community ambiguity, with toy cases indicating that the diagonal model performs particularly poorly in such cases.

3.3 Summary of Two-Community Tests

The results of the two-community aligned and offset tests provide strong evidence for Conjecture 1 and initial evidence for Conjecture 2. In each test, both models consistently identify well-defined and robust global communities. As p approaches q and as the offset value increases, both models perform increasingly poorly. Additionally, the weighted and directed offset test provides initial evidence of a "strong layer community" phenomenon, in which layer communities containing heavily weighted edges can significantly shape global partitioning, especially in the diagonal model. The fact that the non-diagonal model is less sensitive to "strong layer communities" is likely related to its treatment of inter-layer connections, and the resulting emphasis on global structure over that of each individual layer. We expect this to mitigate the effect of individual outlier edges, instead emphasizing the broader global communities.

3.4 Extended Offset Test

The next step in our toy case examination is to repeat the offset case searching for more than two communities. The purpose of this test is to examine how each model handles the ambiguous nodes when the modularity maximization algorithm is allowed to run to completion (i.e. create further subdivisions until $\Delta Q \leq 0$). In order to examine the results of this test in a systematic way, we define a list of 18 logical partitions of the overlapping and non-overlapping regions. We define these partitions as "logical" because they encompass every possible partition of the four offset ambiguous regions that in some way groups together overlapping nodes. We consider this list to be exhaustive, and as such treat all other partitions as illogical.

All parameters are set at the same values as in the previous offset test, i.e. two layer networks with q={0,0.05,…,0.5} and ℓ={1,2,…,25}. The partition created for each value of q and ℓ is compared to all 18 classes of partitions, and recorded under the class to which it belongs. This classification is done by splitting the nodes into the following partition and examining each set separately:

$$\textit{Ambiguous 1}: \{ v_i^1 : i \in [1, \ell] \}$$
$$\textit{Overlapping 1}: \{ v_i^1 : i \in [\ell + 1, 50] \}$$
$$\textit{Ambiguous 2}: \{ v_i^1 : i \in [51, 50 + \ell] \}$$
$$\textit{Overlapping 2}: \{ v_i^1 : i \in [51 + \ell, 100] \}$$
$$\textit{Ambiguous 3}: \{ v_i^2 : i \in [1, \ell] \}$$
$$\textit{Overlapping 3}: \{ v_i^2 : i \in [\ell + 1, 50] \}$$
$$\textit{Ambiguous 4}: \{ v_i^2 : i \in [51, 50 + \ell] \}$$
$$\textit{Overlapping 4}: \{ v_i^2 : i \in [51 + \ell, 100] \}$$

Next, each region is assigned a community tag corresponding to the most frequent community of the nodes within the region, and the region tags are compared to those of each of the 18 partition classes. A full definition of each partition classification and region visualization can be found in Appendix Section 6.1. If a partition does not match any of the

18 classes, it is classified as "Illogical" and recorded as such. Additionally, we slightly tweak the definition of the *Mistake Ratio* (*r')* for these tests:

$$r' = \frac{overlapping\ nodes\ whose\ community\ differs\ from\ region\ tag}{total\ overlapping\ nodes}$$

This definition differs from the previous definition of *r* in that regions are considered on a layer by layer basis, with nodes counted as mistakes only if they are separated from the majority of nodes in their *layer specific* region.

In order to systematically examine the results of this extended version of the offset test, we record the quantities of each class of partition across all trials, and the *illogical* count broken out by *q* value. Once again, these tests are first performed for the undirected-unweighted case, and then for the generalized weighted and directed case. These results are presented in Figures 3.5 and 3.6 below.

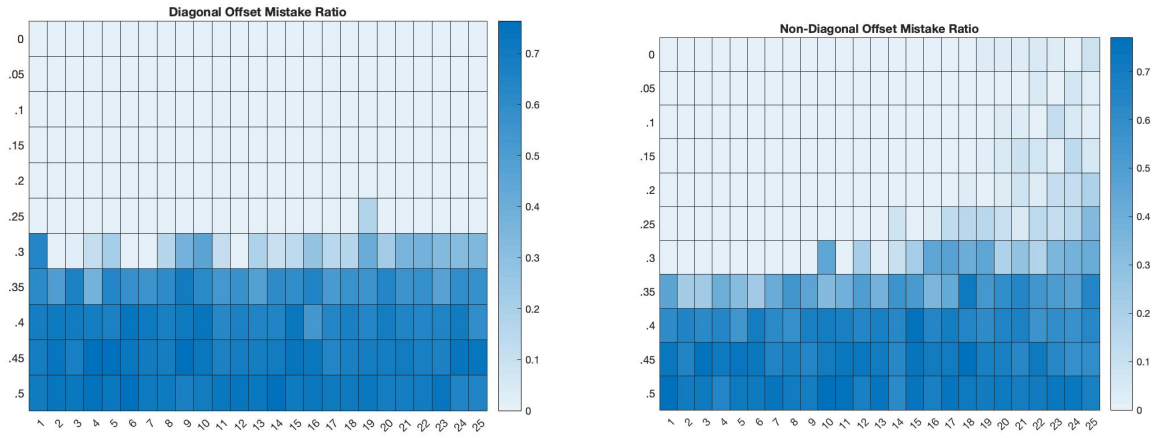

Partition Type	Diagonal Count	Non-diagonal Count
Type 1	146	117
Type 2	0	5
Type 7	0	1
Type 8	1	0
Type 10	0	2
Type 11	0	2
Type 13	0	1
Type 14	0	1
Illogical	128	146

Q Value	Diagonal Illogical Count	Non-diagonal Illogical Count
0	0	5
0.05	1	5
0.1	0	7
0.15	2	9
0.2	2	8
0.25	2	7
0.3	22	15
0.35	24	19
0.4	25	25
0.45	25	23
0.5	25	23

Figure 3.5: Unweighted and undirected results for diagonal (left column) and non-diagonal (right column) models

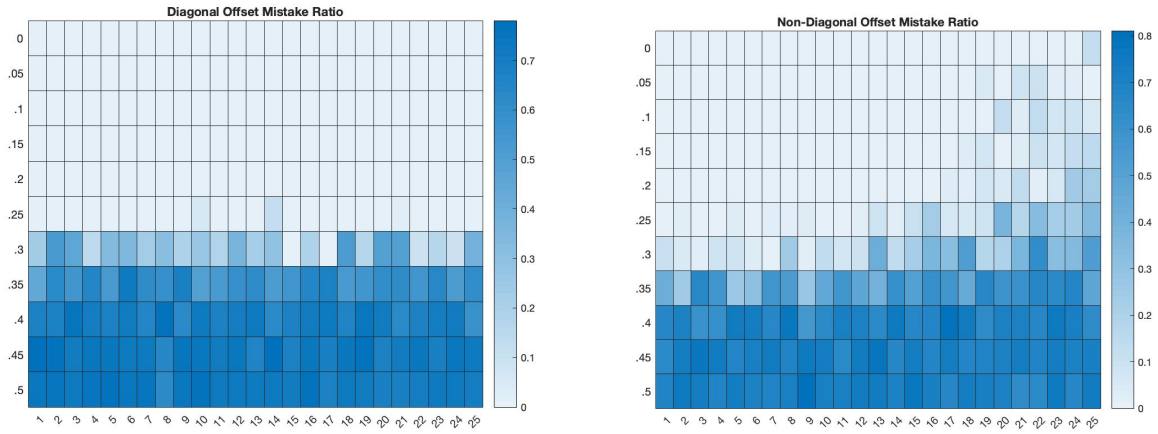

Partition Type	Diagonal Count	Non-diagonal Count
Type 1	145	106
Type 2	0	1
Type 5	4	0
Type 6	1	0
Type 7	0	1
Type 10	0	2
Type 12	0	2
Illogical	125	163

Q Value	Diagonal Illogical Count	Non-diagonal Illogical Count
0	1	8
0.05	1	6
0.1	0	7
0.15	1	7
0.2	1	10
0.25	4	11
0.3	20	19
0.35	24	24
0.4	24	24
0.45	24	24
0.5	25	23

Figure 3.6: Weighted and directed results for diagonal (left column) and non-diagonal (right column) models

3.5 Summary of Extended Offset Test Results

The results of these tests provide additional support for Conjectures 1 and 2. Both models again perform strongly for q values meaningfully less than p, with performance deteriorating significantly as q approaches p. One interesting result here is that the diagonal model shows significantly less sensitivity to the offset value under this modified definition of r'. This indicates that previously occurring mistakes at large offset values were cases of node copies being separated *across layers*, not nodes being separated from their embedded layer community. This represents initial evidence of the tendency of the diagonal model to favor layer communities in cases of large offset values.

Examining the partition classifications provides further support for this preference for layer communities. Both models, and particularly the diagonal model, favor the Type 1 logical partition (over 100 Type 1 partitions in each test), which entails partitioning each layer separately into its embedded communities and then lining up these layer communities such that the overlapping communities are grouped together (see Appendix 6.1 for visualization). The sum of all other logical partitions, which represent all logical global community subdivisions, is at most 12 throughout all of the tests. This indicates that both models, and especially the diagonal model, favor partitioning the nodes by layer communities over partitions that emphasize global structure.

Lastly, breaking out the illogical partitions by q value provides additional evidence that the diagonal model favors layer communities when applied to networks with large offset values. In both the unweighted-undirected and the weighted-directed case, the diagonal model creates very few illogical partitions at low q values. The non-diagonal model, on the other hand, shows sensitivity to the offset value, creating a handful of illogical partitions at low q values (which occur at high offset values). This indicates that for well-defined embedded networks, as the ambiguous region grows larger, the diagonal model continues to emphasize layer communities, separating many node copies in the ambiguous region. The non-diagonal model, on the other hand, begins to generate illogical partitions, most likely due to handling the ambiguous regions in a more nuanced way.

These results indicate that the two models have differing sensitivities to various network characteristics. The diagonal model has a tendency to favor layer communities regardless of the degree of offset, always separating node copies that are in different layer communities, regardless of the quantity of these ambiguous nodes. On the other hand, the non-diagonal model tends to separate "ambiguous" node copies only if they are significantly outnumbered by the "overlapping" nodes, and creates more nuanced (though "illogical") partitions as the ambiguous region grows. These divergent sensitivities illustrate an important difference between the two models, and is another important takeaway from this set of toy case tests.

IV. World Trade Web Analysis

Having examined the performance of the diagonal and non-diagonal multiplex modularity maximization models across well understood toy cases, we now turn to a real-world system, specifically the World Trade Web (WTW). The primary motivation for this application is to understand how each model performs on "real" systems, and to examine the results in the context of the toy case findings. Additionally, this analysis will highlight new avenues for macroeconomic research by applying novel analytical approaches to an important macroeconomic system, specifically the global trade system.

4.1 Data

The WTW is a complex system that has been frequently studied within the field of economics. Early analysis by Serrano and Boguñá (2003), Garlaschelli and Loffredo (2005), and Fagiolo et. al. (2008) studied the structural properties of the single layer WTW network. Serrano et. al. (2006) examined separate layers of the WTW network using a disjoint-layers model. Recently, authors have begun examining the disaggregated WTW as a network on multiple layers, including Barigozzi et. al.'s (2011) dynamic analysis of individual WTW layers, and the initial multiplex WTW analysis of Lee and Goh (2016).

Continued academic interest in the WTW is due to it being a natural application for network analysis, as trade can easily be conceptualized as a network, with countries

represented as nodes, and quantity of trade (in dollars) between countries represented as weighted edges. The application of modularity maximization to such a trade network allows for the identification of sub-communities of highly connected countries, representing trade blocs within the global system. The multiplex modularity maximization models developed in this paper allow for a more nuanced analysis of the WTW system. Specifically, examining the WTW as a multiplex network, with trade disaggregated into commodity components on separate layers, allows for the identification of commodity-specific trading blocs, providing more nuanced insight into global trade patterns.
The identification of highly-connected, commodity-specific trading blocs within the system of global trade can provide valuable macroeconomic and political insight. A high degree of multilateral trade, as identified by the clustering of modularity maximization, indicates economic interdependence. This economic interdependence can be an important input for macroeconomic and foreign policy at the state level. In this sense, our WTW analysis serves to both further the development of multiplex analytical models, and provide novel insights and avenues for future research in the fields of macroeconomics and foreign policy.

4.2 SITC-2 Pairwise Analysis

Throughout this section, we analyze WTW data from 2000 (Feenstra et. al., 2005). In our application, we use the SITC-2 classifications, which disaggregate total trade flows into 66 commodity layers, according to the United Nations' Standard International Trade Classification (SITC) two-digit code, as defined in Appendix Section 6.2. It should be noted that, as a result of disaggregation, many of the SITC-2 commodity layers are quite sparse, with trade between only a handful of nodes. Additionally, one would intuitively expect edge weights to exhibit right skew, with a small handful of heavily weighted edges between the largest interconnected economies. Lastly, there are likely certain highly connected trading blocs that are present on the vast majority of layers, such as the E.U.

We first apply both modularity maximization models to every possible pair of commodity layers. This examination runs parallel to the offset toy case from the previous section. In the toy case, embedded communities were offset by a varying degree across 2 layers, with a high probability of intracommunity edges and a lower (but varying) probability of

intercommunity edges. The results of this analysis were largely supportive of Conjectures 1 and 2. These findings can be translated to the WTW context, with highly similar commodity layers considered as "layers with a low degree of offset (highly overlapping)" and dissimilar commodity layers as "highly offset." We will use this interpretation to develop a series of hypotheses regarding the performance of each model in the SITC-2 pairwise analysis, but first we must define two new metrics:

$$Offset\ Ratio\ (OR) = \frac{mismatched\ nodes}{n}$$

$$Mismatch\ Ratio\ (MR) = \frac{separated\ node\ copies}{n}$$

In the above definition, *mismatched nodes* are defined based on separate, single-layer modularity maximization of each of the two layers of interest. Each single-layer community formed on Layer 1 is matched with a corresponding Layer 2 community to which the majority of its nodes belong, and *mismatched nodes* are counted as those whose community on Layer 2 is different than the community matched with its Layer 1 community. This can be interpreted as the conceptual parallel of ℓ, and one would expect that if OR were calculated for the toy cases, it would be that $OR \approx \frac{\ell}{50}$. As seen in Figure 4.1 below, in which the OR value is calculated for each $\ell \epsilon [1,25]$ on two layer networks as defined in the Offset toy case Section 3.2, this does hold to be true. Thus, *Offset Ratio* represents a meaningful equivalent of ℓ from the toy cases and can be interpreted as such.

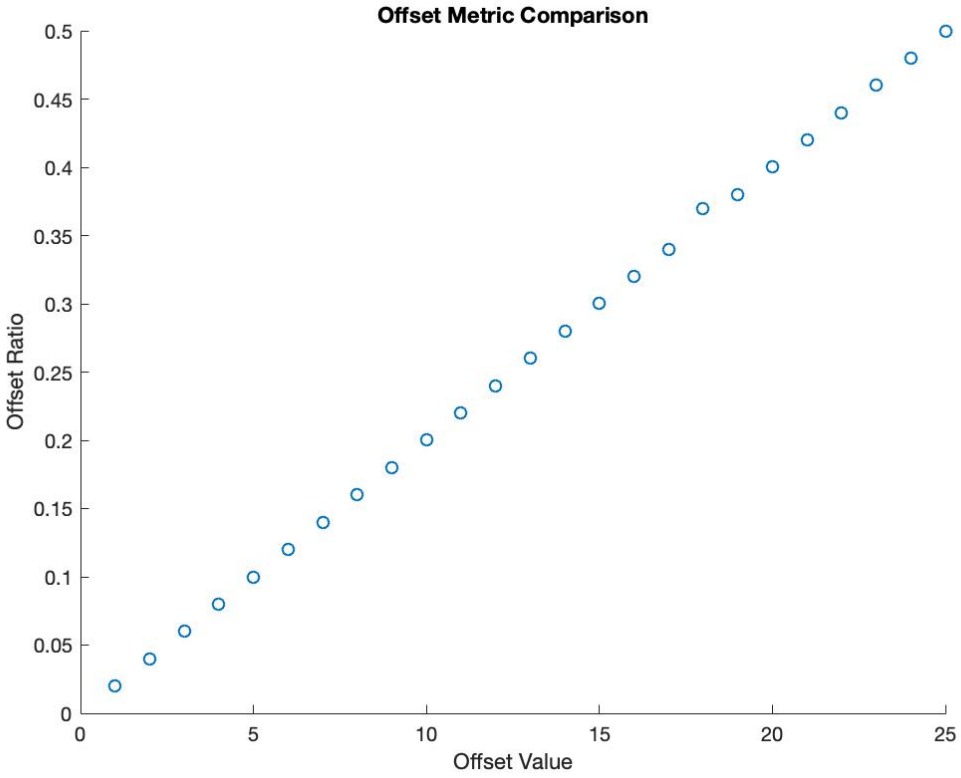

Figure 4.1: Offset Ratio calculated for each value of ℓ from toy case analysis, evaluated at q=0.2 (all other values set as in Section 3.2)

In the definition of *MR, separated node copies* is simply defined as the number of node copies that are separated (i.e. placed in different communities) under a given multiplex modularity maximization model.

Interpreting Conjectures 1 and 2 and the toy case findings in the context of this SITC-2 pairwise test and these newly defined metrics, we hypothesize:

> *Hypothesis 1:* When applied to layer pairs with low *Offset Ratios*, both modularity maximization models will identify the robust global communities, and as a result separate few node copies and produce a low *Mismatch Ratio*. This directly follows from Conjecture 1.

Hypothesis 2: When applied to layer pairs with larger *Offset Ratios*, both modularity maximization models will increasingly favor layer communities, splitting node copies that are offset across layers, resulting in a higher *Mismatch Ratio*. This directly follows from Conjecture 2.

Hypothesis 3: When applied to layer pairs with exceptionally high *Offset Ratios*, performance of the two models will diverge, with the diagonal model strongly favoring layer communities and the non-diagonal model producing more nuanced, though at times illogical, partitions. As this occurred for $\ell > 20$ in the toy cases, we expect it to occur for $OR > \frac{20}{50} = 0.4$ in these trials.

In the following trials, we will test these hypotheses, which together imply a positive relationship between *Offset Ratio* and *Mismatch Ratio* and divergent model performance at extreme *Offset Ratios.* To this end, we present the results of all possible pairings of WTW SITC-2 commodity layers below in Figure 4.2.

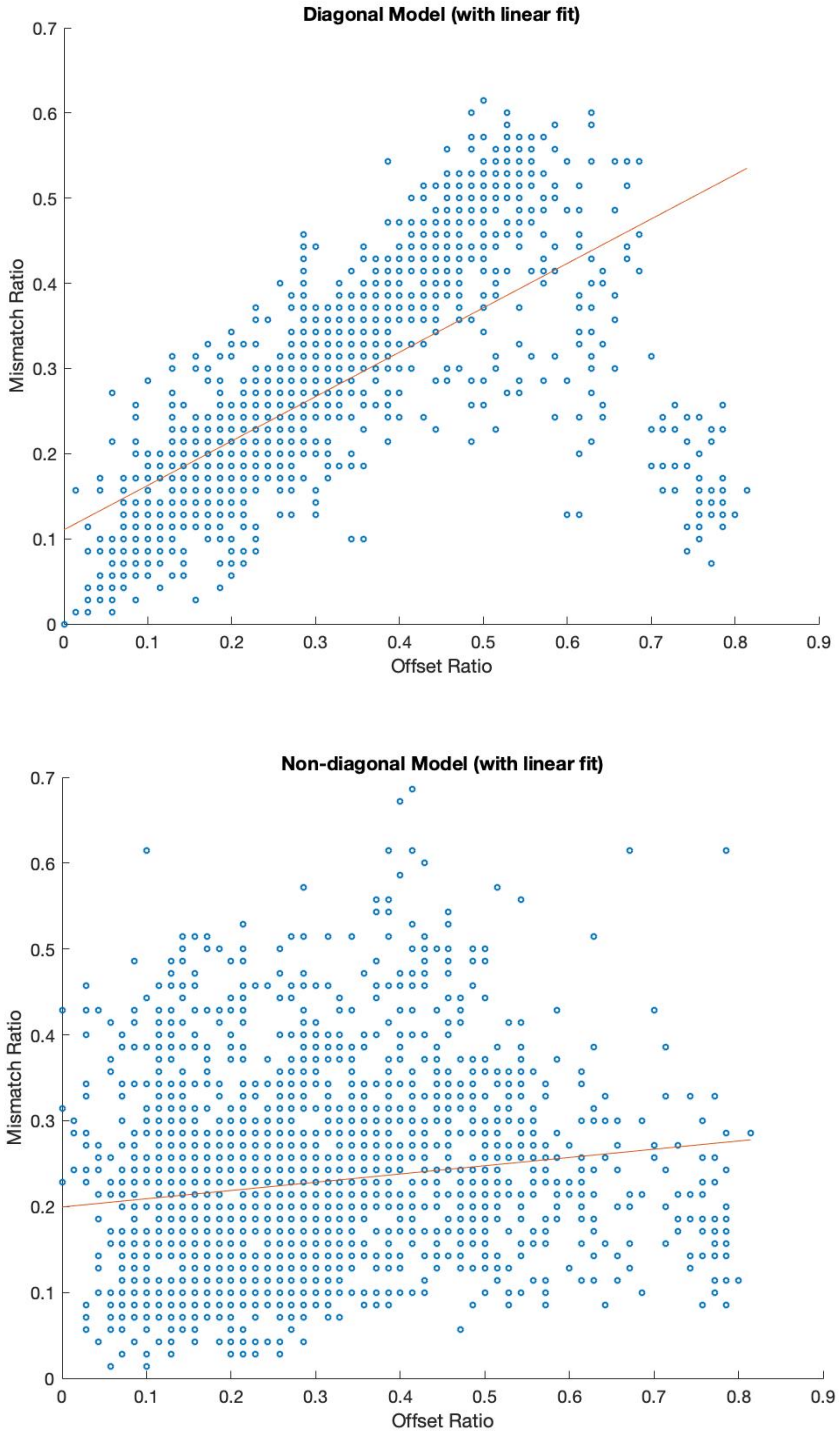

Figure 4.2: Diagonal (top) and non-diagonal (bottom) models' MR as a function of underlying layers' OR

Diagonal: $\widehat{MR} = 0.11 + 0.52OR$ *with* $R^2 = 0.45$

Non-Diagonal: $\widehat{MR} = 0.20 + 0.10OR$ *with* $R^2 = 0.02$

The results from the diagonal model support Hypotheses 1 and 2. There is a clear and strong linear relationship between *OR* and *MR,* with an R^2 value of 0.45, indicating that the diagonal modularity maximization increasingly favors layer communities as the degree of offset increases. This closely mirrors the findings from the offset tests. However, the results do not support Hypothesis 3, as the relationship breaks down at higher offset values, indicating an anomaly not explained by a continued emphasis on layer communities.

Upon close examination of this region of the findings, we hypothesize that this is due to pairings involving layers with sparse edges. Pairings involving a sparse layer as "Layer 2" result in very high *Offset Ratios,* as the sparse layer is split into many communities by single layer modularity maximization, the majority of which contain a single, isolated node. As a result, the Layer 1 communities do not coherently match with these many sparse layer communities, resulting in a large *OR* value. However, such pairings tend to result in very low *Mismatch Ratios* as the communities on the non-sparse layer tend to dominate, forcing node copies together as the sparse layer communities are ignored. We will test this hypothesis, concretely stated below, in the following section.

> *Hypothesis 4:* Pairings involving sparse layers can result in high *Offset Ratios* and low *Mismatch Ratios*, distorting the relationship between the two metrics.

The results from the non-diagonal model do not support our hypotheses. There is no discernable relationship between *Offset Ratio* and *Mismatch Ratio*, with an R^2 value of 0.02, contradicting Hypotheses 1 and 2. Furthermore, this phenomenon is not limited to high offset values, thus contradicting Hypothesis 3. Interpreting these results in light of the performance of the diagonal model, we hypothesize that this is a result of the way that the non-diagonal model handles interlayer connections. Specifically, the non-diagonal model captures more nuanced global communities, and decrease the emphasis on node copy connections. Perhaps this greater analytical complexity explains the decreased strength of the relationship between *OR* and *MR*, as the *MR* is increasingly a product of more subtle network characteristics, and less a product of node copy offset. To put this plainly, perhaps

OR does not fully capture the "different-ness" of the layer pairs, and thus does not strongly explain *MR*. This hypothesis is stated below, and explored further in Section 4.5.

> *Hypothesis 5:* The non-diagonal model's treatment of interlayer connections renders the *Offset Ratio* as insufficient in capturing the differences between layer pairs, and mutes the relationship between *OR* and *MR.*

4.3 SITC-1 Pairwise Analysis

In order to explore (and limit) the effects of sparse layers, we next perform pairwise analysis on SITC-1 WTW data. This represents a somewhat more aggregated version of the trade data, with total flows separated into 10 (instead of 66) commodity layers. See Appendix Section 6.2 for further description of the SITC-1 classifications. By examining broader commodity classes, we ameliorate some of the effects of sparse layers. Additionally, reducing the total possible pairings allows us to better understand commodity pair relationships, and predict the degree of similarity that we would expect within a given pairing.

To contextualize the analysis in this way, we classify all pairings of SITC-1 layers into the following buckets:

> *Supply Chain*: Commodity pairs that are inputs to a common final product, with trade expected to be in part determined by this shared supply chain.
>
> *Similar Consumables*: Commodity pairs that are expected to have end-markets with highly overlapping demographics, and thus similar trade patterns.
>
> *Seemingly Unrelated*: Commodity pairs that do not fall under either of the previous two classifications.
>
> *Sparse*: Commodity pairs involving "Other Commodities" (layer 9), as this is by far the sparsest layer, with only 8.7% of all possible edges taking a non-zero value.

Semi-Sparse: Commodity pairs involving "Beverages and tobacco" (layer 1) or "Animal and vegetable oils" (layer 4), as these are the second and third sparsest layers, with 26.8% and 35.9% non-zero edges respectively.

We perform two versions of the following SITC-1 analysis, one with *Semi-Sparse* pairings classified as they would be (within the first three buckets), and one including the separate *Semi-Sparse* classification. A legend documenting the classification of each layer pairing can be found in Appendix section 6.3. The results of both versions of this aggregated analysis, broken out by pairing type, are shown below.

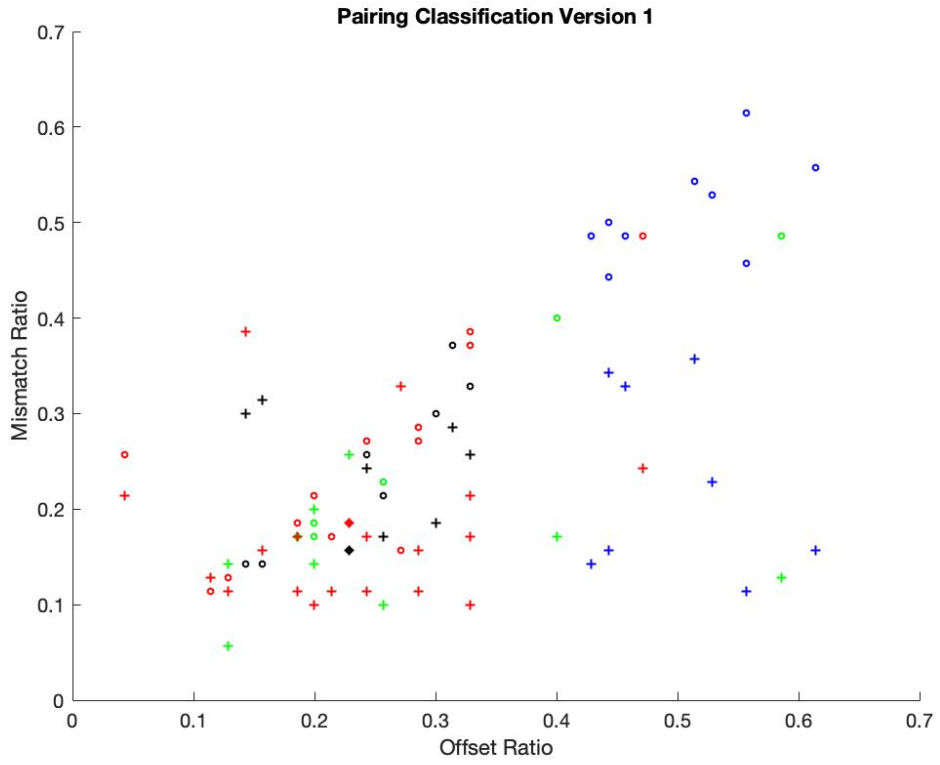

Figure 4.3: Version 1 of the SITC-1 pairwise analysis. Diagonal results shown with circles, non-diagonal shown with +'s. Supply Chain pairs in green, Similar Consumables in red, Seemingly Unrelated in Black, and Sparse in Blue

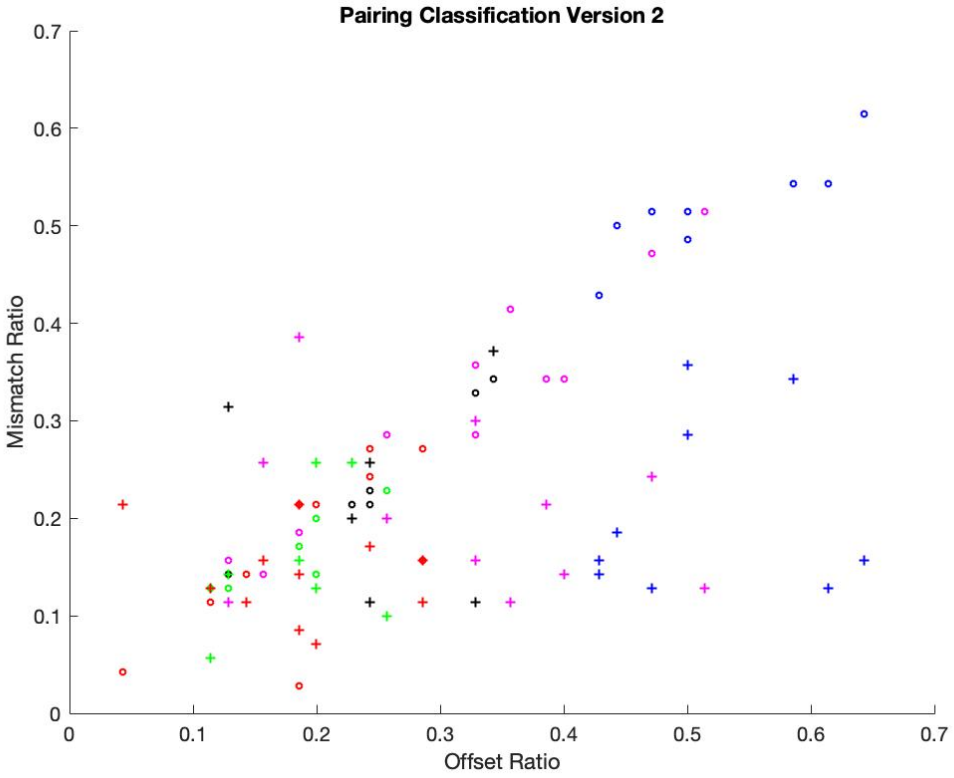

Figure 4.4: Version 2 of the SITC-1 pairwise analysis. Diagonal results shown with circles, non-diagonal shown with +'s. Supply Chain pairs in green, Similar Consumables in red, Seemingly Unrelated in Black, Semi-Sparse in magenta, and Sparse in Blue

The results of this SITC-1 analysis closely support Hypothesis 4. This can be seen in Figure 4.3, in which the *Sparse* pairings are seen to have the largest *Offset Ratios*, and the most varied *Mismatch Ratio* outcomes. Figure 4.4 further supports this hypothesis, as the *Semi-Sparse* pairings are primarily located between all other non-sparse pairs and the *Sparse* pairs in regard to *OR*. Taken together, these findings illustrate the confounding effects of sparse layers, which we will eliminate in the next section by excluding all such layers.

4.4 Non-Sparse SITC-1 Analysis

We next seek to isolate non-sparse pairs and focus on the relationship between *Offset Ratio* and *Mismatch Ratio* for each model when analyzing these pairings. In order to do this, we

formulate a third version of the SITC-1 pairing classifications, excluding all *Sparse* and *Semi-Sparse* pairings. We perform a linear regression of *MR* results on corresponding *OR* values for this third version, and plot the linear fit with results broken out by model and pair type below:

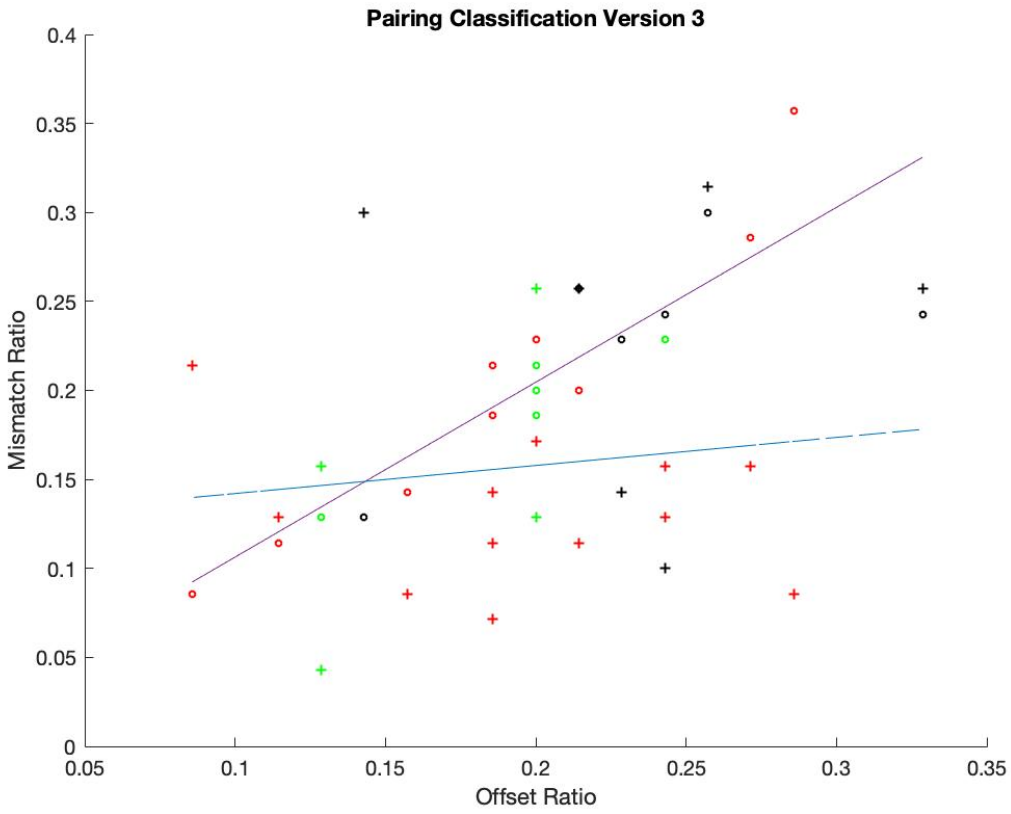

Figure 4.5: Version 3 of the SITC-1 pairwise analysis. Diagonal results shown with circles, non-diagonal shown with +'s. Supply Chain pairs in green, Similar Consumables in red, and Seemingly Unrelated in Black

Diagonal: $\widehat{MR} = 0.01 + 0.98OR$ *with* $R^2 = 0.79$

Non-Diagonal: $\widehat{MR} = 0.13 + 0.16OR$ *with* $R^2 = 0.01$

Examining the results, shown in Figure 4.5, we find strong support for Hypotheses 1, 2, and 4. Comparing diagonal performance in this test versus the SITC-2 analysis illustrates the confounding effects of sparse layers. Previously, the diagonal model yielded an R^2 value of 0.45. However, with sparse pairs excluded, the R^2 value increases to 0.79. Thus, we

conclude that when operating on non-sparse pairs, there is a strong relationship (with a slope that is approximately equal to 1) between the layer pairs' *Offset Ratio* and the *Mismatch Ratio* produced by the diagonal model.

Another interesting result in Figure 4.5 is that the *Seemingly Unrelated* pairs tend to have unusually high *MR* values, both in absolute terms and compared to the corresponding *OR* value. The finding that these pairs have large *MR* values in absolute terms is not surprising as these pairs were identified as those least likely to have common trade patterns, i.e. we expect less similarity across layers and less robust global communities. Thus, by Hypothesis 2, we would expect both modularity maximization models to favor layer communities, separate many nodes, and generate large *MR* values.

However, it is surprising to note that when applied to *Seemingly Unrelated* pairs, the non-diagonal model in particular yields *MR* values that are significantly higher than what would be predicted given the *OR* value. This can be seen in the fact that four of the six non-diagonal *Seemingly Unrelated* results lie above the trend line, and three of these four represent the three highest *MR* values produced by the non-diagonal model, and have very large positive residuals in the linear fit. Interpreting this in the context of Hypothesis 3, we next return to the question of whether *Offset Ratio* is perhaps insufficient in capturing the full extent of dissimilarity across layers that shapes modularity maximization performance, especially by the non-diagonal model that has been shown to create highly nuanced global partitions. We address this question in the next section by building a new metric for "different-ness."

4.5 Ego Differences Analysis

Our *Offset Ratio* metric clearly captures some of the differences in layer structure within pairings, but it is not fully satisfactory. Specifically, *OR* does not seem to fully explain why the non-diagonal model produces particularly large *Mismatch Ratios* when applied to *Seemingly Unrelated* pairings. We hypothesize that this may stem from our definition of *OR.* Specifically, our *OR* algorithm matches each Layer 1 community with the Layer 2 community to which the majority of its nodes belong. This allows for trading blocks of large

countries that trade heavily across almost all commodities (such as the United States and Mexico, or European Union countries) to drown out more nuanced differences in trading patterns of smaller countries. In an attempt to better capture these differences, we next turn to ego networks. For background, see Newman (2010, 44-46).

Each node *v* in a single layer network *N* has a corresponding ego network *G*, which is defined as follows:

$$V(G) = \{v\} \cup \{N(v)\}; \quad E(G) = \{(u,v): (u,v) \epsilon E(N)\} \cup \{(u,w): u,w \epsilon N(v) and\ (u,w) \epsilon E(N)\}$$

In other words, the ego network corresponding to node *v* contains all of *v*'s neighbors, the corresponding edges, and any edges between neighbors. Comparing the ego networks of node copies across two layers allows for another measure of "different-ness."

Specifically, we define a metric called *Ego Difference (ED)*. Note in the following definitions that G_1 refers to node *v's* ego network on Layer 1 and G_2 refers to its ego network on Layer 2:

$$v's\ ratio\ of\ shared\ neighbors = RN(v) = \frac{|(V(G_1) \setminus \{v\}) \cap (V(G_2) \setminus \{v\})|}{\max\left(|(V(G_1) \setminus \{v\})|, |(V(G_2) \setminus \{v\})|\right)}$$

$$v's\ ratio\ of\ shared\ 2nd\ order\ edges = RE(n) = \frac{|(E(G_1 - v) \cap E(G_2 - v)|}{\max\left(|E(G_1 - v)|, |E(G_2 - v)|\right)}$$

$$ED = 1 - 0.5\left(\frac{\sum_{v \epsilon V(N)} RN(v) + RE(v)}{n}\right)$$

Defined as above, our *Ego Differences* metric ranges from 0 (the two layers are identical, every node has the exact same ego network on each layer) to 1 (the two layers do not share a single edge in common, every node has distinct ego networks on the two layers with no shared neighbors). It is important to note here that *ED* completely ignores edge weights. The metric captures shared edges across layer pairings, but does so in a binary way. This

aspect of the formulation is an intentional attempt to downplay the confounding effects of tight knit trading blocks of large countries, as was discussed in the previous section.

With the preceding definitions in mind, Figure 4.6 below shows a graph of *MR* vs. *ED* for both models, broken out by pairing classification version 3:

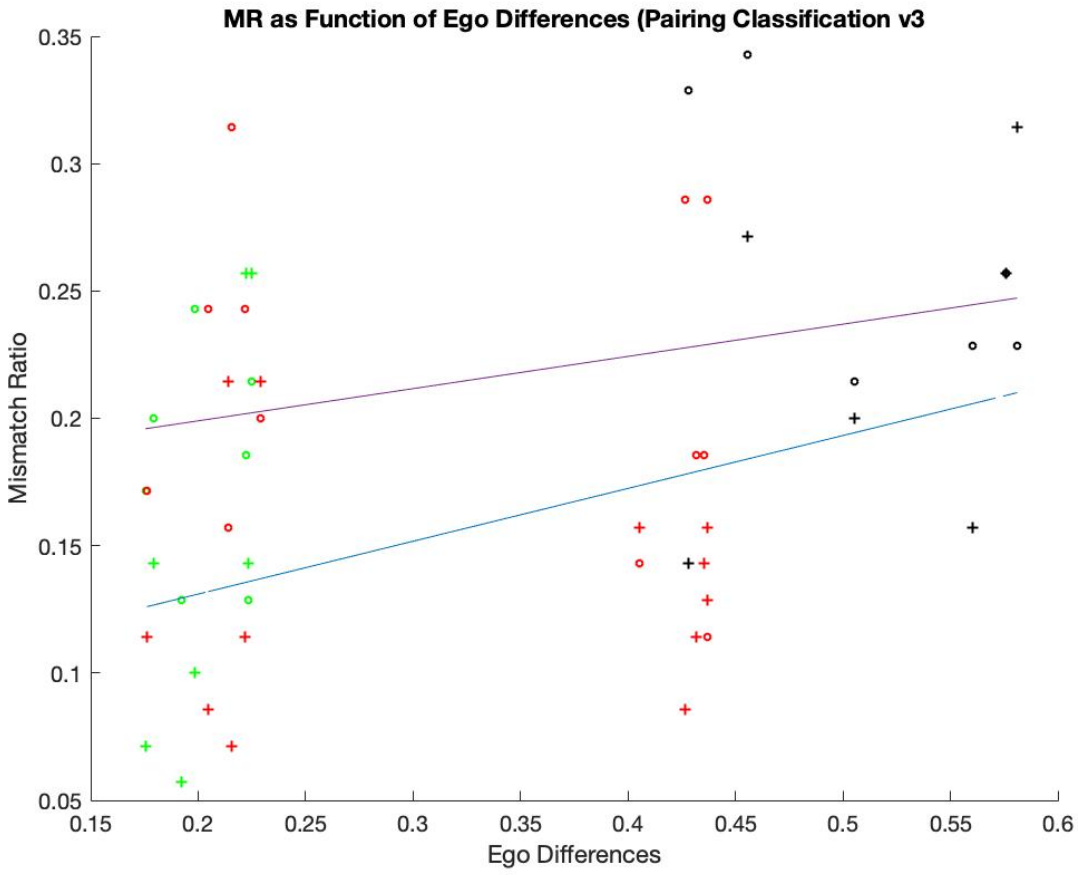

Figure 4.6: Diagonal (circles) and non-diagonal (+'s) models' MR as a function of underlying layers' ED for non-sparse layer pairs

Diagonal: $\widehat{MR} = 0.17 + 0.13ED$ *with* $R^2 = 0.08$

Non-Diagonal: $\widehat{MR} = 0.09 + 0.21ED$ *with* $R^2 = 0.18$

The above results are very interesting, especially when compared to Figure 4.5. The first thing to note is that the *Ego Differences* metrics does a very good job of capturing the

differences between layer classifications, with *Seemingly Unrelated* pairs (black) resulting in significantly higher *ED* values than either *Similar Consumables* (red) or *Supply Chain* (green) pairs. This confirms our hypothesis that *Offset Ratio* was not fully capturing "different-ness" within layer pairs, most likely due to the over-powering effects of robust cross-commodity trading blocs such as the EU. In this sense, the *ED* metrics is helpful in teasing out differences between layer pairs.

A second noteworthy result is that the diagonal model produces a stronger relationship between *OR* and *MR* (slope= 0.98 and $R^2 = 0.79$) than between *ED* and *MR* (slope= 0.13 and $R^2 = 0.08$), whereas the non-diagonal model produces a stronger relationship between *ED* and *MR* (slope= 0.21 and $R^2 = 0.18$) than between *OR* and *MR* (slope= 0.16 and $R^2 = 0.01$).

This offers strong support for Hypothesis 5 and provides the basis for Conjecture 3, as stated in the Introduction. The diagonal and non-diagonal models' divergent sensitivity to varying types of cross-layer "different-ness" represents an important analytical distinction. The global partition produced by the diagonal model is primarily a function of the network's *Offset Ratio,* or put more simply, a function of the degree to which the layer communities broadly line up when considering weighted edges. On the other hand, non-diagonal model partitions are primarily a function of the network's *Ego Differences,* i.e. first and second order similarity of node-neighborhoods with edges treated as binary. This makes sense in the context of our previous findings, as we expect the non-diagonal model's handling of interlayer edges to amplify more nuanced second order similarities/differences between layers, while we expect the diagonal model to be shaped by layer communities containing highly weighted edges. This finding, concretely captured in Conjecture 3, provides important insight and helps inform analytical decisions made by future multiplex modularity maximization practitioners.

V. Conclusion

Modularity maximization is a powerful tool for analyzing network structure and identifying communities of highly connected nodes. These tools are especially valuable to those examining highly complex networks, as the community detection resulting from modularity maximization facilitates network visualization and highlights important structural characteristics that would otherwise be obscured within the complex system. For this reason, the extension of modularity maximization models to multiplex networks, which are significantly more complex and unintelligible than single layer networks, represents a particularly important development within the field of network analysis. By examining the performance of two multiplex modularity maximization models across different families of networks, and building a set of theoretically and empirically justified conjectures, we help to arm future researchers with valuable community detection tools and inform sound multiplex network analysis.

Our first main finding is that both the diagonal and non-diagonal multiplex modularity maximization models successfully identify well-defined and robust global communities. We state this claim explicitly as Conjecture 1, and find strong justification across a series of toy case and World Trade Web tests. In the toy case analysis, we apply both models to a series of networks with embedded layer communities that are aligned across layers, and find that both models partition the nodes as we expect, identifying the robust embedded global communities. Additionally, when applied to toy networks with embedded layer communities that are offset across two layers, both models tend to cluster the node copies that are within overlapping (similar) communities on both layers. Lastly, we find that both models tend to separate few node copies when applied to WTW networks with highly similar layer community structure, indicating that both models correctly identify the robust global communities in these cases. Together, these findings provide strong support for Conjecture 1 and indicate the baseline validity of both models.

Secondly, we find that both models tend to emphasize layer communities when applied to networks with highly different, or offset, layers. We state this claim as Conjecture 2 and

find strong justification across all of the tests we perform. We find initial evidence for this claim in the offset toy case analysis as both models favor Type 1 partition classifications, which represent partitions driven predominantly by layer community structure, over all other "logical" partitions. In particular, the extended offset test indicates that the diagonal model especially emphasizes layer communities in highly offset cases. We find evidence for this in the diagonal model's tendency to separate node copies but rarely separate nodes within embedded layer communities (low *r'* but high *r* values), and the model's continued preference for Type 1 partitions even at extreme offset values for which such a partition separates many node copies. We find strong support for Conjecture 2 in the WTW analysis, as we find a positive relationship between layer "different-ness" and the degree to which node copies are separated across a variety of tests. This indicates that when applied to networks containing layers with highly different community structure, both models emphasize layer communities and as a result separate many node copies. Together, these tests provide strong justification for Conjecture 2.

Lastly, we find that the two models are sensitive to different types of network layer "different-ness." This finding, stated concretely in Conjecture 3, represents perhaps the most valuable insight of our analysis. As mentioned above, we find a positive relationship between layer "different-ness" and node copy separation in the partitions produced by both models across all WTW tests. However, the two models are sensitive to two different conceptualizations of "different-ness." The diagonal model produces a strong relationship between *Offset Ratio,* which primarily captures the extent to which heavily weighted layer communities are aligned across layers, and *Mismatch Ratio,* which captures the extent to which node copies are separated. However, the non-diagonal model shows essentially no relationship between these two metrics. Instead, the non-diagonal model produces a stronger relationship between *ego differences,* which measures the similarity of node copy neighborhoods across layers, and *Mismatch Ratio* (two metrics that are mainly unrelated for the diagonal model). This finding is very valuable in informing future research, allowing practitioners to choose the proper modularity maximization model based on the type of community that they seek to identify.

By building out and justifying each of these three conjectures, we hope to open new avenues for research within the field of network analysis. Most proximally, our World Trade Web analysis can serve as a jumping off point for further macroeconomic research. Our systematic examination of each models' performance when applied to all possible layer pairs provides thorough evidence of the tendencies of each model when applied to the WTW system. Armed with this understanding, future researchers can better analyze the global trade system and identify commodity-specific trading blocs. Informed by Conjecture 3, these researchers can fine tune their model choice and specifications to emphasize the type of layer similarity that they consider to be most important. The results of such analysis would likely be highly valuable to the makers of macroeconomic and foreign policy.

Furthermore, the value of our findings is not limited to the WTW network. Multiplex networks are a natural representation for a multitude of complex systems. Research on such systems has been broadly hampered by the lack of well understood and conceptually justified multiplex analytical tools. Our three conjectures illustrate the baseline validity, analytical tendencies, and differing sensitivities of two valuable multiplex modularity maximization models. As a result, future researchers can apply these models with a strong understanding of their performance, and an informed ability to tailor model choice and specifications to their analytical goals. Thus, we see our findings as the equivalent of the user's manual for a power tool. Such a power tool can be incredibly useful, but its value goes only as far as the user's understanding of how to use it. By providing a series of well justified conjectures on model performance, we hope to arm future researchers with the understanding to reap the rewards of powerful multiplex modularity maximization models.

VI. Appendix

6.1 Offset Toy Case Partition Classifications

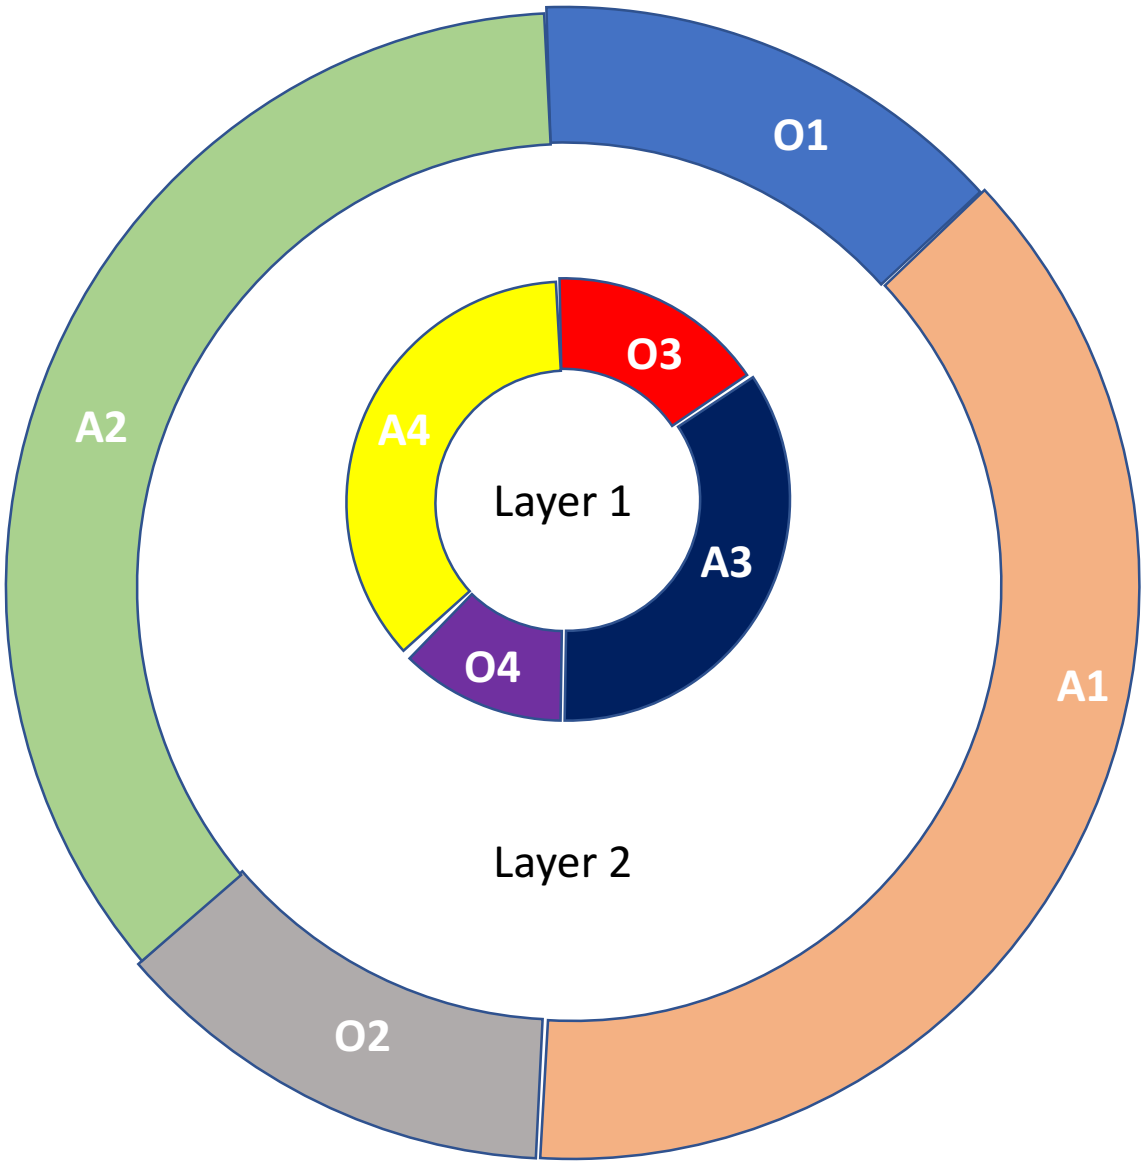

Figure 6.1: Stylized layer community regions (labeled in white text)

$$\textit{Ambiguous 1}: \{\, v_i^1 : i \in [1, \ell] \,\}$$
$$\textit{Overlapping 1}: \{\, v_i^1 : i \in [\ell + 1, 50] \,\}$$
$$\textit{Ambiguous 2}: \{\, v_i^1 : i \in [51, 50 + \ell] \,\}$$
$$\textit{Overlapping 2}: \{\, v_i^1 : i \in [51 + \ell, 100] \,\}$$

Ambiguous 3: $\{ v_i^2 : i \in [1, \ell] \}$

Overlapping 3: $\{ v_i^2 : i \in [\ell + 1, 50] \}$

Ambiguous 4: $\{ v_i^2 : i \in [51, 50 + \ell] \}$

Overlapping 4: $\{ v_i^2 : i \in [51 + \ell, 100] \}$

PARTITION TYPE	**A1**	**O1**	**A2**	**O2**	**A3**	**O3**	**A4**	**O4**
TYPE 1	i	i	ii	ii	ii	i	i	ii
TYPE 2	i	ii	iii	iv	i	ii	iii	iv
TYPE 3	i	i	ii	ii	i	i	ii	ii
TYPE 4	i	ii	ii	i	i	ii	ii	i
TYPE 5	i	i	ii	ii	iii	iv	iv	iii
TYPE 6	i	ii	iii	iv	v	ii	vi	iv
TYPE 7	i	i	ii	iii	ii	i	i	iii
TYPE 8	i	ii	iii	iii	iii	ii	i	iii
TYPE 9	i	ii	iii	iv	iii	ii	i	iv
TYPE 10	i	ii	iii	iv	v	ii	iii	iv
TYPE 11	i	ii	iii	iv	i	ii	v	iv
TYPE 12	i	ii	i	i	i	ii	i	i
TYPE 13	i	i	i	ii	i	i	i	ii
TYPE 14	i	i	i	ii	iii	i	i	ii
TYPE 15	i	i	ii	iii	i	i	i	iii
TYPE 16	i	ii	i	i	i	ii	iii	i
TYPE 17	i	ii	iii	iii	iii	ii	iii	iii
TYPE 18	i	i	i	i	i	i	i	i
ILLOGICAL	Any partition that doesn't match any of the above types							

Table 6.1: Region groupings by partition type, region tags shown in roman numerals

6.2: SITC Code Descriptions

SITC 1 Digit Codes

Code	Commodity
0	Food and live animals
1	Beverages and tobacco
2	Crude materials,inedible,except fuels
3	Mineral fuels etc
4	Animal and vegetable oils and fats
5	Chemicals and related products,n.e.s.
6	Basic manufactures
7	Machinery,transport equipment
8	Miscellaneous manufactured articles
9	Goods not classified elsewhere

Table: 6.2: SITC 1 Digit code descriptions

SITC 2 Digit Codes

Code	Commodity
00	Live animals
01	Meat and meat preparations
02	Dairy products and birds' eggs
03	Fish and fish preparations
04	Cereals and cereal preparations
05	Vegetables and fruit
06	Sugars,sugar preparations and honey
07	Coffee,tea,cocoa,spices
08	Feeding stuff for animals
09	Miscellaneous edible products and preparations
11	Beverages
12	Tobacco and tobacco manufactures
21	Hides,skins,furskins,raw
22	Oil seeds,oleaginous fruits
23	Crude rubber (incl.synthetic)
24	Cork and wood
25	Pulp and waste paper
26	Textile fibres and their wastes
27	Crude fertilizers and crude minerals
28	Metalliferous ores and metal scrap
29	Crude animal,vegetable materials n.e.s.
32	Coal,coke and briquettes
33	Petroleum and products
34	Gas,natural and manufactured
35	Electric current
41	Animal oils and fats
42	Fixed vegetable fats and oils
43	Processed animal or vegetable oils,etc.
51	Organic chemicals
52	Inorganic chemicals
53	Dyeing,tanning and colouring material
54	Medicinal and pharmaceutical products

55	Perfume,cleaning etc.preparations
56	Fertilizers,manufactured
57	Plastics in primary forms
58	Plastics in non-primary forms
59	Chemical materials and products,n.e.s.
61	Leather,dressed fur,etc.
62	Rubber manufactures,n.e.s.
63	Wood and cork manufactures
64	Paper,paperboard and articles thereof
65	Textile yarn,fabrics,made up articles,etc.
66	Non-metallic mineral manufactures,n.e.s.
67	Iron and steel
68	Non-ferrous metals
69	Manufactures of metals,n.e.s.
71	Power generating machinery and equipment
72	Machinery for specialized industries
73	Metal working machinery
74	General industrial machinery n.e.s.
75	Office machines and adp machines
76	Telecommunications and sound recording equipm
77	Electric machinery,n.e.s.and parts
78	Road vehicles
79	Other transport equipment
81	Prefabr.buildings;sanitary,lighting etc.fixtrs
82	Furniture and parts thereof
83	Travel goods,handbags and sim.containers
84	Articles of apparel and clothing accessories
85	Footwear
87	Instruments and apparates n.e.s.
88	Photographic equipment,optical goods etc.
89	Miscellaneous manufactured articles,n.e.s.
91	Postal packages not classified according to kind
93	Special transactions and commodities not classified
96	Coin (not gold coin or legal)
97	Gold,non-monetary

Table: 6.3: SITC 2 Digit code descriptions

6.3 Pairwise Commodity Type Classifications

0-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9
0-2	1-3	2-4	3-5	4-6	5-7	6-8	7-9	
0-3	1-4	2-5	3-6	4-7	5-8	6-9		
0-4	1-5	2-6	3-7	4-8	5-9			
0-5	1-6	2-7	3-8	4-9				
0-6	1-7	2-8	3-9					
0-7	1-8	2-9						
0-8	1-9							
0-9								

Table 6.4: Pairwise layer classifications version 1, i-j refers to the two-layer test examining SITC-1 layers i and j. Supply Chain pairs in green, Similar Consumables in red, Seemingly Unrelated in Gray, and Sparse in teal

0-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9
0-2	1-3	2-4	3-5	4-6	5-7	6-8	7-9	
0-3	1-4	2-5	3-6	4-7	5-8	6-9		
0-4	1-5	2-6	3-7	4-8	5-9			
0-5	1-6	2-7	3-8	4-9				
0-6	1-7	2-8	3-9					
0-7	1-8	2-9						
0-8	1-9							
0-9								

Table 6.5: Pairwise layer classifications version 2, i-j refers to the two-layer test examining SITC-1 layers i and j. Supply Chain pairs in green, Similar Consumables in red, Seemingly Unrelated in Gray, Semi-Sparse in magenta, and Sparse in teal

References

Banerjee, Abhijit, et al. "The Diffusion of Microfinance." *Science*, vol. 341, no. 6144, 26 July 2013, doi:10.1126/science.1236498.

Barigozzi, Matteo, et al. "Community Structure in the Multi-Network of International Trade." *Communications in Computer and Information Science Complex Networks*, 2011, pp. 163–175., doi:10.1007/978-3-642-25501-4_17.

Bazzi, Marya, et al. "Generative Benchmark Models for Mesoscale Structure in Multilayer Networks." *Cornell University,* 27 Nov. 2016, https://arxiv.org/abs/1608.06196.

DeFord, Daryl R. "Matched Products and Dynamical Models for Multiplex Networks." *Dartmouth College*, 2018.

DeFord, Daryl R., and Scott D. Pauls. "A New Framework for Dynamical Models on Multiplex Networks." *Journal of Complex Networks*, 2017.

Fagiolo, Giorgio, et al. "On the Topological Properties of the World Trade Web: A Weighted Network Analysis." *Physica A: Statistical Mechanics and Its Applications*, vol. 387, no. 15, 15 June 2008, pp. 3868–3873., doi:10.1016/j.physa.2008.01.050.

Fagiolo, Giorgio, et al. "The Evolution of the World Trade Web: A Weighted-Network Analysis." *Journal of Evolutionary Economics*, vol. 20, no. 4, 2009, pp. 479–514., doi:10.1007/s00191-009-0160-x.

Feenstra, Robert, et al. "World Trade Flows: 1962-2000." *NBER Working Paper*, vol. 11040, Jan. 2005, doi:10.3386/w11040.

Garlaschelli, Diego, and Maria I. Loffredo. "Structure and Evolution of the World Trade Network." *Physica A: Statistical Mechanics and Its Applications*, vol. 355, no. 1, 1 Sept. 2005, pp. 138–144., doi:10.1016/j.physa.2005.02.075.

Jenelius, Erik, et al. “Importance and Exposure in Road Network Vulnerability Analysis.” *Transportation Research Part A: Policy and Practice*, vol. 40, no. 7, Aug. 2006, pp. 537–560., doi:10.1016/j.tra.2005.11.003.

Lee, Kyu-Min, and K.-I. Goh. “Strength of Weak Layers in Cascading Failures on Multiplex Networks: Case of the International Trade Network.” *Scientific Reports*, vol. 6, no. 1, 23 May 2016, doi:10.1038/srep26346.

Menichetti, Giulia, et al. “Weighted Multiplex Networks.” *PLoS ONE*, vol. 9, no. 6, 6 June 2014, doi:10.1371/journal.pone.0097857.

Newman, Mark E. J. *Networks: An Introduction*. Oxford University Press, 2010.

Serrano, M. Ángeles, et al. “Correlations in Weighted Networks.” *Physical Review E*, vol. 74, no. 5, 13 Nov. 2006, doi:10.1103/physreve.74.055101.

Serrano, Ma Ángeles, and Marián Boguñá. “Topology of the World Trade Web.” *Physical Review E*, vol. 68, no. 1, 11 July 2003, doi:10.1103/physreve.68.015101.

Stanley, Natalie, et al. “Clustering Network Layers with the Strata Multilayer Stochastic Block Model.” *IEEE Transactions on Network Science and Engineering*, vol. 3, no. 2, 25 Mar. 2016, pp. 95–105., doi:10.1109/tnse.2016.2537545.