NETWORKS AND INTERNATIONAL RELATIONS: THE MEASUREMENT OF ALLIANCE PORTFOLIO SIMILARITY

Peter M. Li
Department of Political Science
University of California, Los Angeles
li@ucla.edu

June 2, 2005

Abstract

This paper employs the concepts and methods of social network theory and cluster analysis to improve the measurement of alliance portfolio similarity. Analysts have used this measure which compares nations' lists of military allies to ascertain the degree to which nations have common security interests. While this would seem to be a straightforward exercise, we have undermined our efforts in two ways. First, we have ignored the contribution of “distant” or transitive alliance relations. This oversight has lead to both an over- and an underestimation of alliance portfolio similarity. Second, we have overlooked the fact that when comparing portfolios, there are actually two valid notions of similarity. One gives credit to both the common presence and the common absence of alliances. The other gives credit only to their common presence. For the case at hand, we have chosen the former when the latter would be more appropriate.

This paper is part of a larger project which seeks to show how network theory (e.g., Wasserman and Faust 1994; Newman 2003) and cluster analysis (e.g., Kaufman and Rousseeuw, 1990) can benefit the study of international relations. While metaphors about networks and clusters abound in our writings, we have yet to really leverage those ideas to the degree that others have done in math, sociology and physics. This is unfortunate because a network perspective – the recognition that events and/or actors are linked together and that that resulting structure of that interaction can have an important role in understanding and explaining politics - is an to which, I would imagine, many of us are sympathetic. More importantly, this lacunae is important because networks are a unique source of "whole is greater than the sum of the parts" effects, one that differs from usual suspects of the distribution of power (Waltz 1959, 1979) or the distribution of information (e.g., Fearon 1994, Schultz 1999). In fact, some would even go as far as to argue that the network perspective is *the* starting point to understanding any social behavior. (e.g., Burt 1992) The basis for such a claim begins with the argument that existing images of interaction are either too coarse or too fine. On one hand, there are those who view international relations as the product of the dynamics within a single, unitary international system. As such, these analysts effectively dismiss the role of "functionally equivalent" component states. On the other hand, there are those who view international relations as the cumulative product of the pair-wise interaction of states. These analysts ignore or overlook the contribution made by secondary, tertiary and higher order intermediaries. While the former position runs the risk of throwing away information by aggregating away differences, the latter runs the risk of giving too much credit to what may simply be idiosyncrasies. In contrast, a network perspective argues that interaction is neither as global nor as local as standard perspectives would have us believe. In fact, a network perspective argues that the modal form of interaction involves an intermediate number of heterogeneous actors whose opportunities, choices and information are a function of both an actor's position in the network and of overall shape, size and structure of that network.

While this view is admittedly a somewhat more complex view of interaction than we are accustomed to, I would argue that the core intuition about the importance of the network of relations and ties between actors and events deserves more attention than we have heretofore given. Thus, this paper aims to demonstrate the importance of a network perspective by showing how even the most basic, bare-bones notion of a network, one based on transitive links, can aid us in the study of international relations. To that end, this paper has two objectives. First, I present a solution the somewhat arcane problem of measuring the similarity of military alliance portfolios (i.e., nations' list of alliance partners). I have chosen this problem because it clearly and intuitively shows how a network perspective can produce a more accurate measure of alliance portfolios similarity than is possible with current perspective or methods. Second, I use this paper as an opportunity to show how one can do network analysis. In particular, I show how one can use cluster analysis to calculate the similarity of alliance portfolios. By adopting this pragmatic, application-specific approach, I hope demonstrate why we need to be more familiar with, if not fluent in, the concepts and methods of network analysis.

I CONCEPTUAL ISSUES

Alliance portfolios

Analysts have used the similarity of nations' respective alliance portfolios, the list of a nation's military allies, as a measure of the similarity or commonality of those nations' underlying security interests. (Bueno de Mesquita, 1975) Simply put, the closer the lists are, in terms of overlapping membership, the closer the common security interests are said to be. (Altfeld and Bueno de Mesquita, 1979; Morrow 2000) When studying the effect of alliances, the alliance portfolio approach has an advantage over the more common approach of simply looking at whether or not pairs of states are allied (i.e., dummy variables). By looking at the entire set of alliance relations rather than just that nation's specific relationship with individual allies, the alliance portfolio approach incorporates more information about a nation's alliance behavior. Thus, when we then cross-reference one nation's alliance portfolio against that of another's, we are better able to distinguish instances in which an alliance is the result of a "real" commonality of security interests and when they are simply temporary "marriages of convenience."[1]

Besides its intuitive appeal, the alliance portfolios method benefits from the fact that its calculation appears to be a simple, straightforward exercise. This was pretty much the case until Signorino and Ritter (1999) argued that the specific measure employed by Bueno de Mesquita (1975), the nonparametric *tau-b* statistic, suffers from conceptual and empirical problems. At the core of Signorino and Ritter's criticism is that a "correlational" coefficient, like *tau-b*, are measures of association which only describes how closely two variables move together when what we really want is a measure of similarity which would involve a item-by-item comparison of the list of alliance partners. In its place, Signorino and Ritter (1999) offer a spatial[2] measure, *S,* which addresses three deficiencies they found when using Bueno de Mesquita's method: first, a negative association may not be indicative of dissimilarity; second, some identical portfolios cannot be measured with association; third, sometimes association may change but similarity stays the same. (Signorino and Ritter 1999, 121-22). Today, in addition to the simple dummy variable, *S* has become something of a standard measures in the field of quantitative studies of international relations. This is demonstrated by its inclusion in EUGene, the popular data management application. (Bennett and Stam 2000) That said, I would nevertheless argue that *S* does not completely solve the problems that Signorino and Ritter set out to address. In fact, in some

[1] Work in a similar spirit but which use different methodologies are exemplified by Sirovich's use of singular value decomposition to analyze the Rehnquist Supreme Court (Sirovich, 2003) and Poole and Rosenthal's (1985) use of NOMINATE scores to analyze of legislative roll calls.

[2] By "spatial," Signorino and Ritter mean the following: "We assume that a state makes choices over a number of policy dimensions; and that the vector of its multiple policy choices - i.e., its revealed policy portfolio - represents a point in (foreign) policy space. Based on this, our conception of 'similarity' is very specific: the closer two states are in the policy space - i.e. the closer their revealed policy positions - the more 'similar' their revealed policy positions. The further apart two states are in the policy space, the more dissimilar their revealed policy preferences." (Signorino and Ritter 1999, 126) In other words, Signorino and Ritter project each nation's alliance portfolio into Euclidean space and interpret similarity as the distance between the various portfolios within that space. In contrast, the notion of network similarity I develop below argues that while we can use the similarity of alliance portfolios to identify clusters and networks of alliances, we cannot compute the similarity or distance between mutually exclusive networks (i.e., binary space). In essence, each individual alliance network occupies its own separate space rather than residing together in a single common space.

ways, *S* falls short of capturing the notion of alliance portfolio similarity that both they and Bueno de Mesquita describe.

The problem is partly methodological but is really more an issue of perspective. To put it simply, we have made two errors in the way have constructed our measures of alliance portfolio similarity. First, we have overlooked the contribution that transitive linkages between alliances can have on any measure of similarity. In particular, we have undercounted the contribution made by “distant” alliance relations which can lead to both an over- and an underestimation of alliance portfolio similarity. This is what I call the *mapping problem*. Second, we have overlooked the fact that when it comes to doing an item-by-item comparison of lists, there are actually two valid notions of similarity. One gives credit to *both* the common presence and the common absence of an alliance (i.e., symmetric). The other gives credit to *either* the common presence or the common absence of an alliance (i.e., asymmetric). For the case of alliance portfolios, we have chosen symmetric measures when an asymmetric measure would be more appropriate. This is what I call the *counting problem*.

The mapping problem

The first problem with existing methods is that they approach the measurement of alliance portfolio similarity solely from the perspective of individual nations. While it is obvious that we need to calculate portfolio similarity from the perspective of individual nations (or pair thereof), what we have overlooked is that in order to accurately measure alliance portfolio similarity between *nations* we first need to identify and map out the relationships between *alliances*. Skipping this intermediate step would be like trying to estimate the driving distance between two cities without the benefit of a road map. While the distance “as the crow flies” may be a good estimate; ultimately, it is the shape of the network of roads between the two cities that determines the actual driving distance let alone whether we can even get there. Analogously, I argue that it is the map of relations among alliances that ultimately determines the degree of the similarity between nations’ alliance portfolios.

To demonstrate why the alliance map is important, consider how existing methods deal with the aphorism that "a friend of a friend is a friend" (i.e., transitive relations). Let's say we have a simple network of three nations transitively linked via two pair-wise alliances (Figure 1). If *A* is allied with *B*, and *B* is allied with *C*, then existing measures will report that *A* and *C* have identical alliance portfolios since they both share a single common ally. However, as the degree of separation between alliances increases, existing methods run into a problem. By simply extending the chain by one additional link we get the following result: if the *A* is allied with *B*, *B* is allied with *C*, and *C* is allied with *D*, existing measures will report that the *A* and *D* will have zero alliance portfolio similarity (Figure 2). While one might argue that such "distant" relations are somehow less relevant, it should nevertheless be self-evident that such a situation is not the same as when there is no link whatsoever between *A* and *D* (Figure 3).

What we need is a way to give partial credit to "distant relations" within an alliance network (Figure 2) while reserving the category of zero similarity to those alliances that are truly unrelated – those that are members of mutually exclusive alliance networks (Figure 3). This is where a map of alliances comes into play. An alliance network map identifies all related alliances, however distant that relationship, but also identifies all unrelated alliances. In fact, this ability to map alliance networks becomes increasingly important as we move from the stylized daisy chain example above to real data. Since real alliances can differ in terms of the number of member states and since individual nations can themselves belong to multiple alliances, the relationships between alliances can quickly become rather complex. In this way, a map can serve as an accounting system to help us keep track of all of those complex relations. Once we have mapped out the networks of alliances, we can go back trace the paths between individual nations to compare their respective alliance portfolios. However, before we can do that there remains one more issue which affects the construction of a measure of alliance portfolio similarity – the very definition of similarity.

Figure 1 Mapping Problem

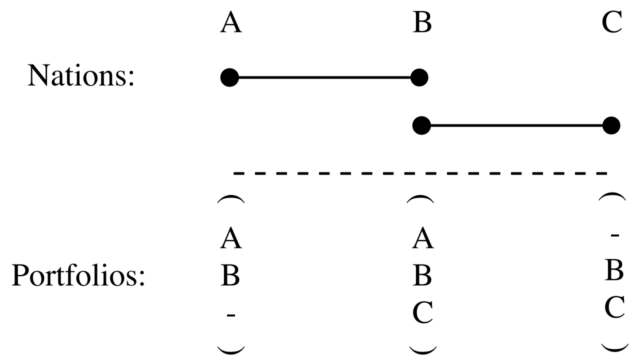

Figure 2 Transitive Link

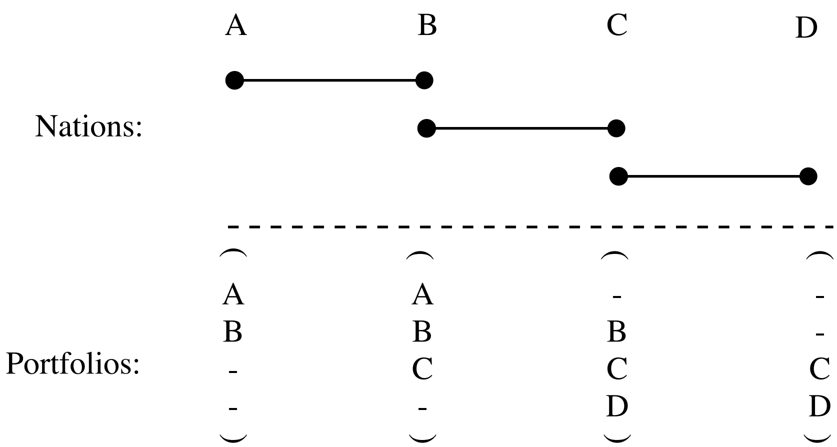

Figure 3 No Transitive Link

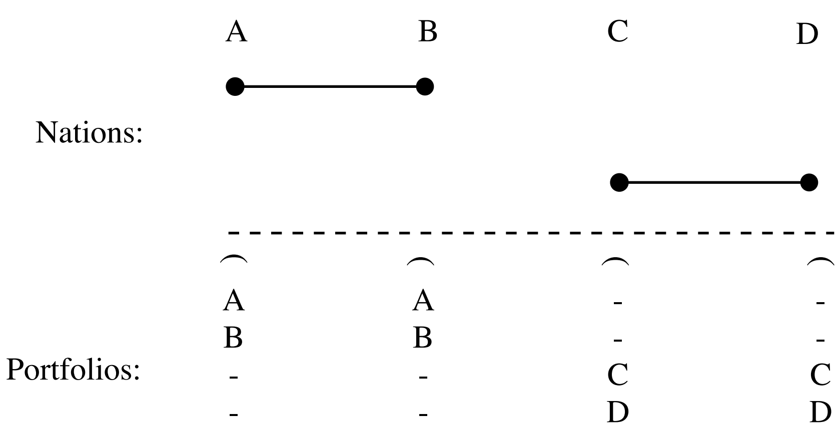

The counting problem

The second problem with existing methods is that we have employed symmetric measures of similarity which gives credit to both the common presence *and* the common absence of alliance ties. In its place, I argue that we should adopt an asymmetric measure that only gives credit for the common presence of alliance ties. This choice is important because even if we were to try to rectify existing methods by adopting the network perspective described above, those modified methods would still not generate the "correct" alliance map. To casual observers, the difference between symmetric and asymmetric measures may seem trivial or bit convenient. Nevertheless, there are substantive conceptual and empirical reasons for the distinction. (Kaufman and Rousseeuw, 1990) With symmetric measures, how we choose to code a variable is simply a convention and does not affect our analysis. For example, when we use a dummy variable to code political parties in a two-party system (e.g. Democrat=1; Republican=0), our analysis is unaffected if we switch coding (i.e. Democrat=0; Republican=1). With an asymmetric measure, this is not the case. While it may be easy to argue that those who are not Democrats are Republicans and vice versa, one cannot as easily make the same distinction for alliances. Those who are not our allies may indeed be our opponents but they may also be neutral states, irrelevant states or simply informal (unobserved) allies.[3] This is why an asymmetric measure is more appropriate. Unlike a symmetric measure that assumes that other nations are either "with us or against us," asymmetric measures take into consideration only what states have in common and leaves the rest aside.

---

[3] Signorino and Ritter (1999) explicitly recognize this problem: "Recall that the inclusion of irrelevant states in alliance portfolios tends to lead to more positive tau-b scores and hence to inflated judgments of the similarity of states' policy positions. S is also affected by the inclusion of irrelevant states: when most of the elements … fall in to the (0, 0) category, S will correctly indicate that two states' alliance portfolios are quite similar in that they tend to have the same type of alliance commitments with each of their partners." (127) Nevertheless, they write: "The obvious solution to this problem would be to ensure that the domain of alliance portfolios is appropriately specified for each empirical application. When used as indicators of foreign policy similarity, states' alliance portfolios should include only those states that can reasonably be considered 'policy relevant'…. Since the alliance data do not allow us to distinguish irrelevant observations of 'no alliance' from strategically meaningful observations of 'no alliance,' narrowing the domain of relevant states by geography or diplomatic activity is probably the best we can do for now."(Signorino and Ritter, 1999, p. 124) However, their solution only shifts the burden from the similarity measure to either some weighting factor or to the more difficult step of identifying which states are politically relevant for a particular state.

There are actually a large number of symmetric and asymmetric measures of similarity and dissimilarity. (Kaufman and Rousseeuw, 1990) Each of these measures differ in how they weigh four component counts. For a measure of alliance portfolio similarity, these component counts are the allies that are common to both nations' alliance portfolios (a), the allies that are unique to a respective nation's portfolio (b and c), and the nations that are not part of either nations' portfolio (d). In terms of the distinction between symmetric and asymmetric measures, (a) represents the common presence of alliance members while (d) represents the common absence. Thus, the basic symmetric measure is the *matching coefficient*, (a + d) ÷ (a + b + c + d), whereas the basic asymmetric measure is the *Jaccard coefficient*, (a) ÷ (a + b + c). The specification of these basic measures help us to further clarify the difference between symmetric and asymmetric measure through the use of a Venn diagram (Figure 4). In some sense, the issue of symmetric versus asymmetric measures boils down to the relative size of (d) and how much credit, if any, should be given to the common absence of alliances. While all nations are potential allies and should be included as entries in an alliance portfolio, the fact that these "zero" observations are not necessarily our opponents means that we should not give credit to the common absence of alliances in our similarity score. This becomes particularly important if alliances are rare events with few participants. In such situations, (d) will be relatively large and the overwhelming majority of entries in a nation's alliance portfolio will be these problematic "zeros."

II NETWORK ANALYSIS

Matrices and graphs

The raw data of network analysis are the relations among and between actors and/or events. This distinction between relations among actors and relations between actors and events plays an important role in our analysis of alliance networks. Matrices which describes relations among actors (i.e., within a single class of objects) are square in that they have the same number of rows and columns. Examples of such data include such things as information on which senators go to lunch together, information about which nations are geographically contiguous or information about the frequency of intermarriage among

Figure 4 Symmetric v. Asymmetric Measures of Similarity

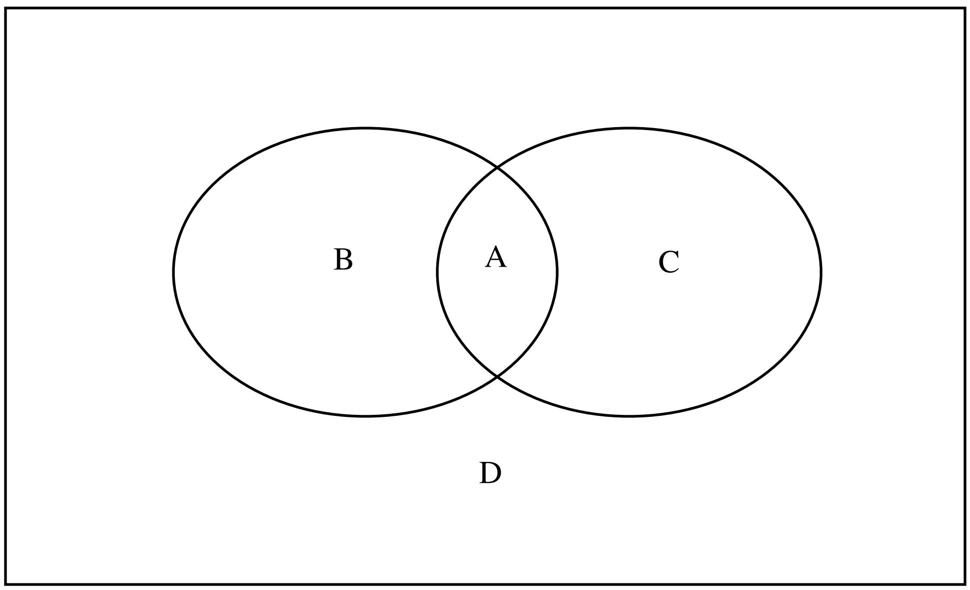

Symmetric (Matching Coefficient) = (A+D)/(A+B+C+D)

Asymmetric (Jaccard Coefficient) = (A)/(A+B+C) different ethnic or racial groups. In the network analysis literature, these data structures are called adjacency matrices. In contrast, matrices which describes relations between actors and events (i.e., between two classes of objects) may not be square since the number of events need not be the same as the number of actors. Examples of this type of data include such things as information on membership on corporate boards (individuals and corporations), legislative roll-calls (legislators and bills), and military alliance membership (nations and alliances). These data structures are called incidence matrices.

To provide perspective on the overall "topography" of relations in a network, analysts illustrate the information contained in both adjacency and incidence matrices as graphs. These network graphs are useful not only because they are inherently asymmetric but also because they provide a complete picture of all direct and indirect (i.e., transitivities) relations among actors and/or events. In Figure 5, we can see a graph of the relations between defensive alliances in the year 1816. The numbers represent a given alliance's Correlates of War numerical code. The graph shows the following. First, we see one disconnected or unrelated alliance, 2010. Such an object is called an isolate. Second, we can see that four alliances, labeled 2005, 2006, 2007 and 2008 all have direct ties with one another. Third, we can see that while alliance 2000 has a direct tie to alliance 2008, it only has indirect ties to 2005, 2006 and 2007. However, since data on military alliances contains information on both alliances and their member states, this representation is only a partial representation of alliance relations since it only looks at the data from the perspective of the 6 alliances and excludes information about the 16 member states. If we shift perspective from the alliance to the member-states, we can get the complementary picture illustrated in Figure 6. While this helps to see the relations among individual states, we lose the information about which alliance particular states belong to. While it may be possible to identify subsets (i.e., alliances) in the graph, the fact that some states can belong to multiple alliances may make such a strategy increasingly difficult especially as the complexity of relations among alliance increases. To address this, network analysts have developed graphs specifically for this type of data. In Figure 7, we have a *bipartite* graph which incorporate information about both alliances, represented as squares, and member states, represented as circles. Here we see that large structure observed in the first two graphs is a product of the

Figure 5 Alliance Perspective - 1816 Defensive Alliances

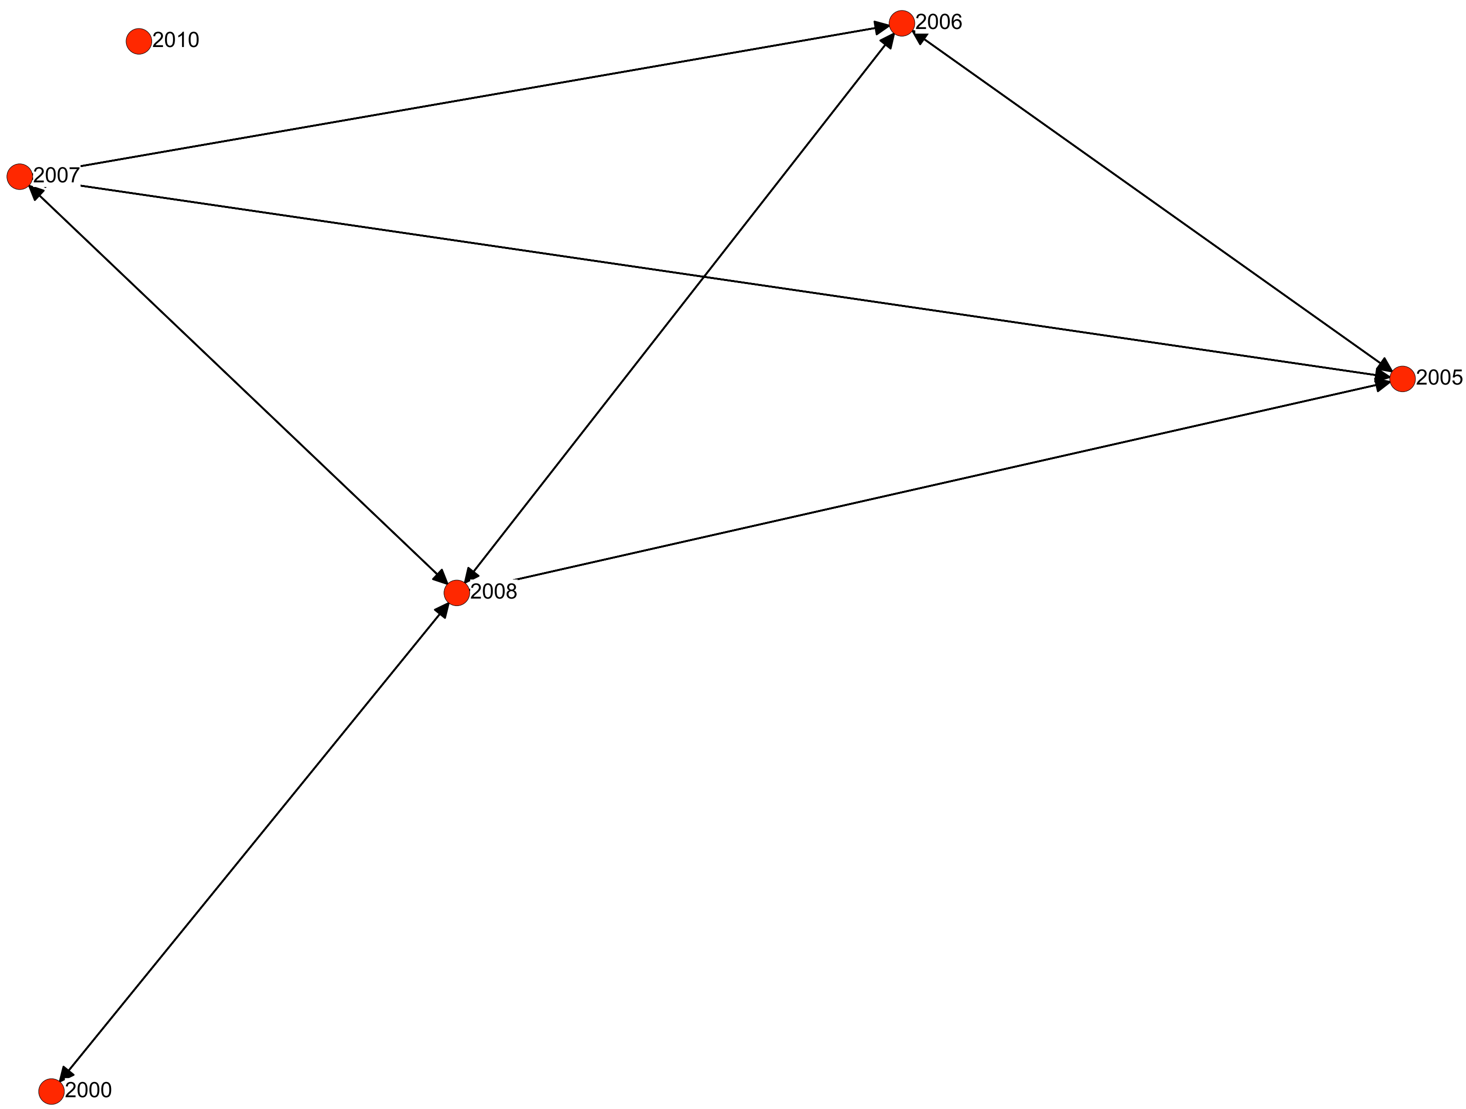

Figure 6 Alliance Member Perspective - 1816 Defensive Alliances

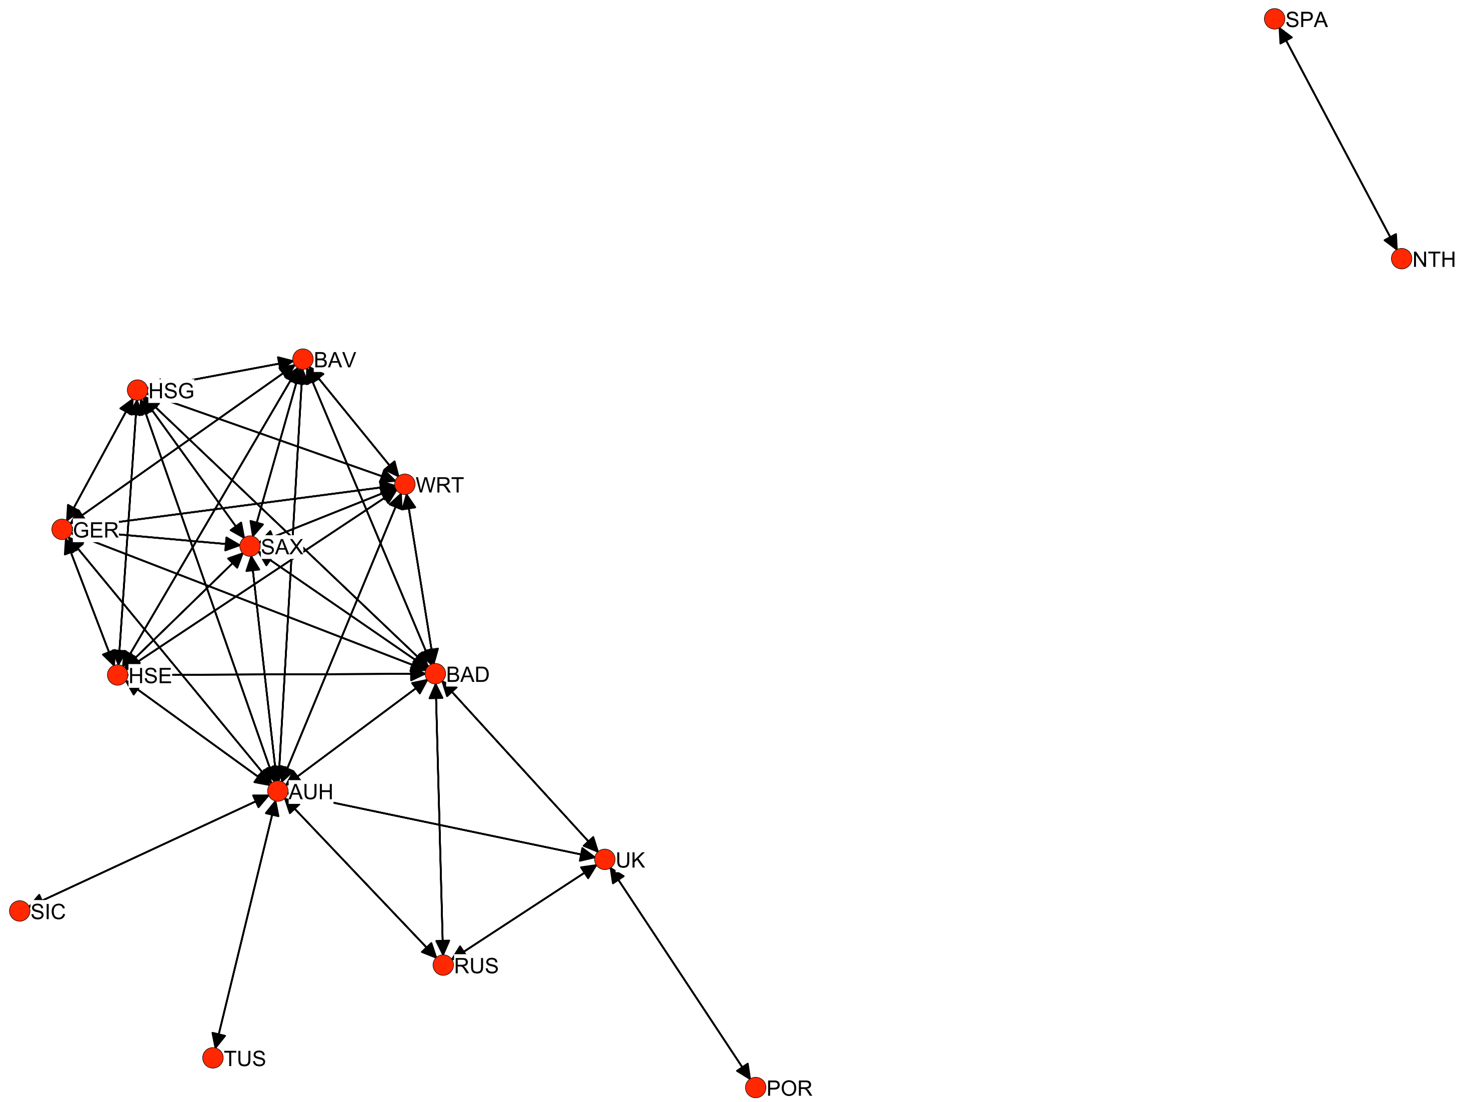

Figure 7 Bipartie Graph - 1816 Defensive Alliances

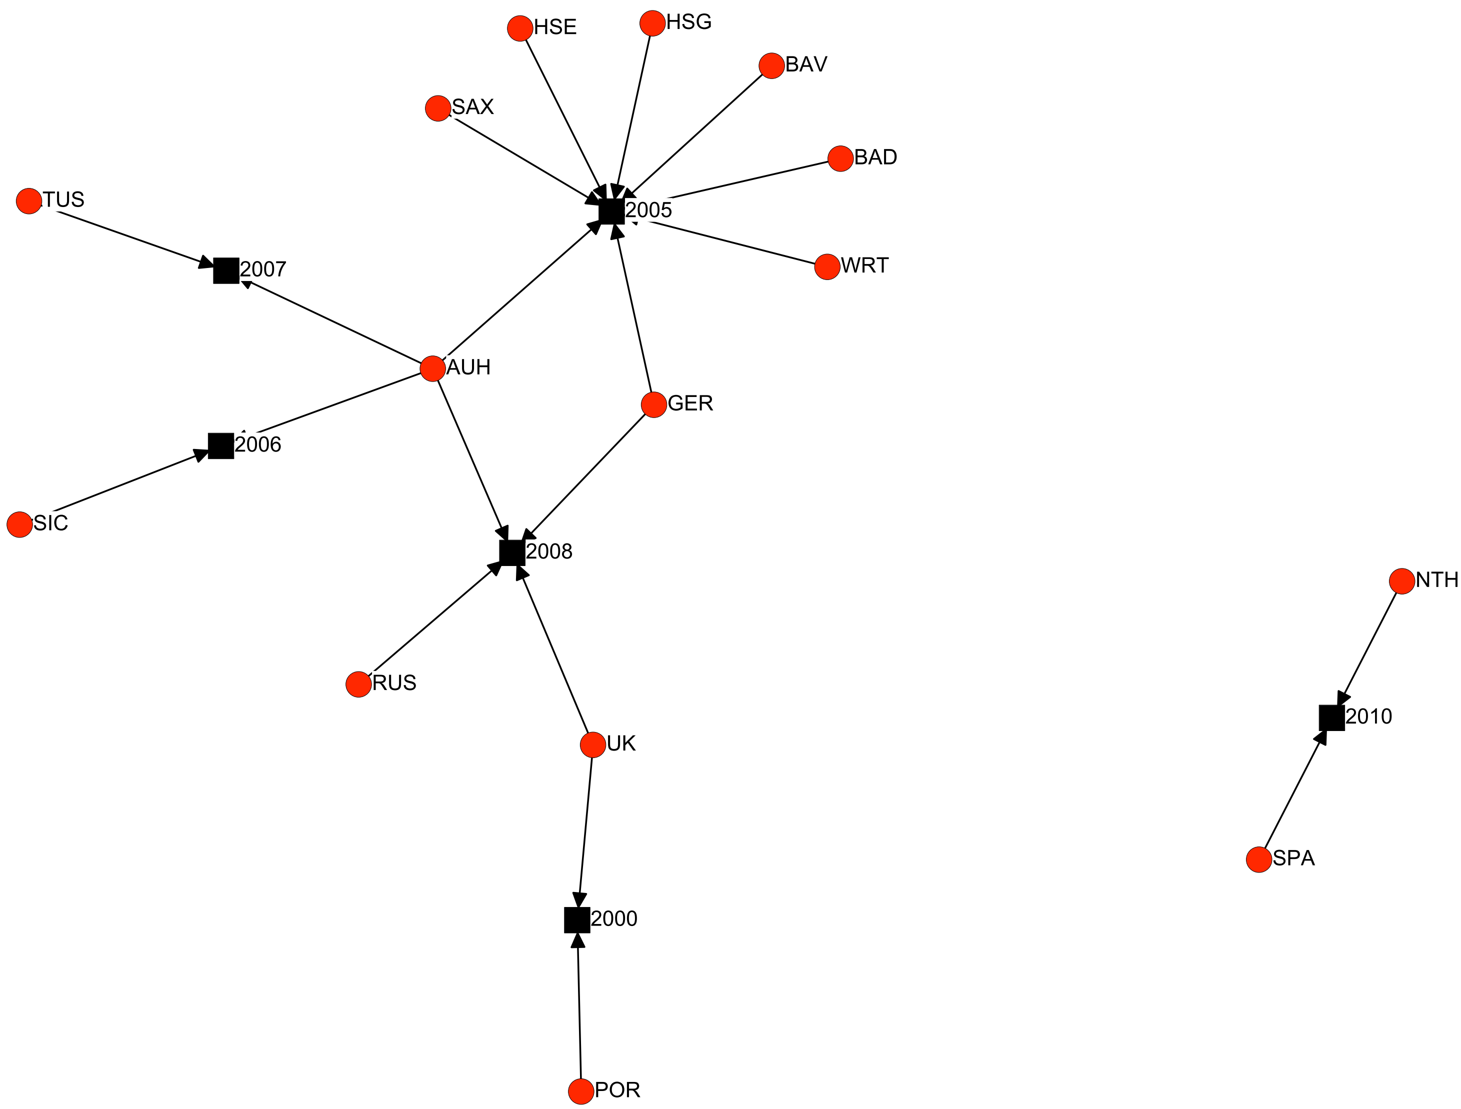

fact that Austria-Hungary (AUH) and Germany/Prussia (GER) are serving as bridges linking together alliances 2005, 2006, 2007 and 2008 into a large, single network.

By literally mapping out the overall structure, the relative position of individual nations and the links and relations which tie nations and alliance together, these graphs provide a nice and succinct demonstration of the value and importance of the network perspective. That said however, these graphs can easily become visually cluttered as the number of alliances and alliance members increase. More importantly, these graphs are not the best way of easily assessing the similarity of nations' alliance portfolios. For both these reasons, it would be useful to use the metric of distance or similarity as a way of organizing and reducing the data. To that end, in the next section, I describe how I use cluster analysis to group together alliances and nations with similar degree of portfolio similarity to create a hierarchical map of alliance relations.

Cluster analysis

There are many flavors of cluster analysis. The basic differences among the various forms are the following. First, cluster analysis can be *agglomerative*, *divisive* or *optimal*. The difference between the first two depends on whether we start with small clusters and works our way up to a single large cluster or whether we break down a single cluster into smaller individual clusters. In contrast, optimal methods tries to find some specified number of clusters in the data. Second, the output of cluster analysis could either be *hierarchical* or *partitional* depending upon whether or not the specific method produces nested classes of increasingly distant or unrelated clusters, or whether the output simply separates observations into different groups or classes. Third, depending upon the number of characteristics used (i.e., univariate v. multivariate), the criteria for identifying clusters can be *monothetic* or *polythetic*. In this paper, I use agglomerative hierarchical clustering based on pooled annual data of three categories of military alliances (i.e. entente, nonaggression and defensive) taken from the Correlates of War project.

Agglomerative hierarchical clustering is an algorithm that works via local rather than global optimization. Based upon some criteria of distance, similarity or relatedness (I use the Jaccard measure of similarity discussed above), the algorithm first ranks those alliances with the greatest and least degree of overlapping membership and then merges the two most similar alliances into one cluster. The algorithm then iteratively compares the similarity of that new cluster with the remaining alliances and/or clusters. To do this, however, the algorithm does not recalculate all the possible distances between the individual elements of clusters. Instead, it uses a measure of inter-cluster similarity. The three most common measures of inter-cluster distance are single-linkage or nearest neighbor (i.e., the shortest distance between clusters), complete-linkage or furthest neighbor (i.e., longest distance between clusters) and average distance. I use a single-linkage or nearest neighbor measure because that allows me to identify all related alliances (i.e., it captures all transitive relations). In fact, single-linkage is sometimes called a "friends of friends" clustering strategy. In any event, the agglomerative hierarchical clustering algorithm merges objects together into larger and larger clusters of alliances starting with those that are most closely related, continuing with those that are most tenuously related, and finally ending with those that are, for all intents, unrelated or mutually exclusive. At this point, we have exhausted the data and are left with a single object and a hierarchical list of distances between alliances.

Dendrograms

This list of hierarchical distance or similarity are typically graphed as a dendrogram which is a kind of "family tree" diagram. In Figure 8, I present the network map for alliances in the year 1816. It presents the information from Figure 5 in hierarchical form based on the degree of similarity between alliances due to the overlapping membership of individual states. The vertically-oriented numbers within the graph represents an alliance's Correlates of War identification code. The bracketed horizontally-oriented abbreviations are the Correlates of War country code of that alliance's member states. Alliances that are most closely related are at the bottom of the graph. Alliances that are the least related are closer to the top. The vertical scale, labeled "height," is the calculated distance between clusters. The height of a given horizontal branch represent the distance between all alliances or clusters at that level. Thus,

Figure 8 Alliance Network

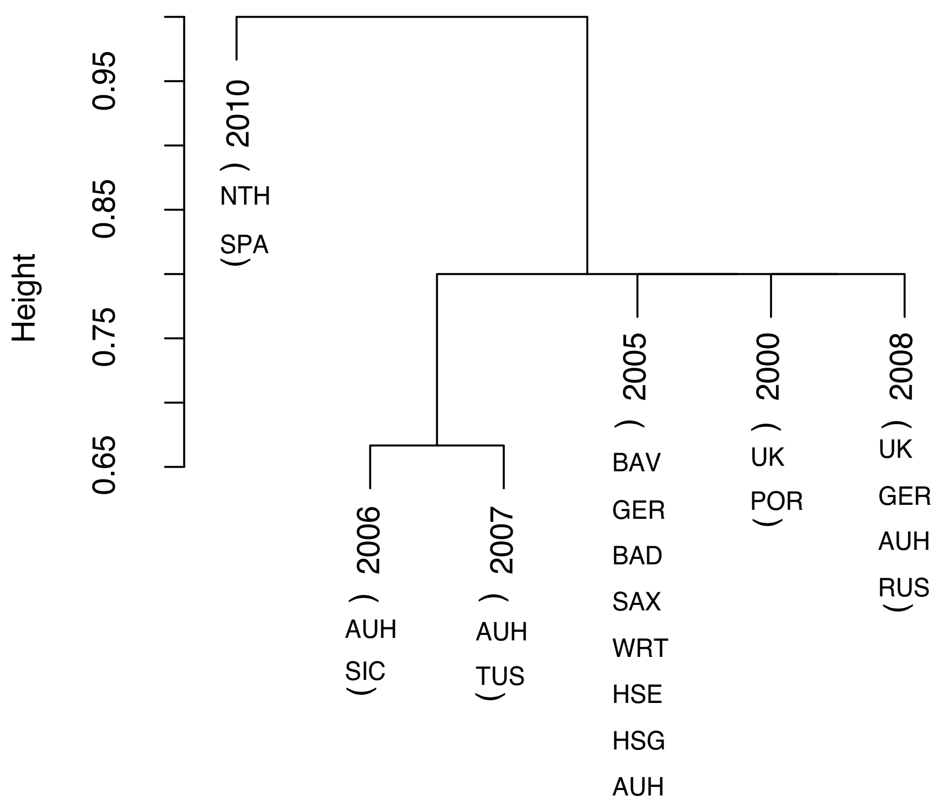

alliances 2000, 2005, 2006, 2007 and 2008 have a *dis*similarity score of approximately 0.8 on a scale from zero to one.

Three things are worth noting when reading a dendrogram. First, you may notice that the scale is reversed from the language used up to this point. This is because the convention is to use dissimilarity (i.e., one minus the similarity score) rather than similarity as the distance metric. In Figure 8, alliances 2000, 2005, 2006, 2007 and 2008 have a dissimilarity score of 0.8 but a similarity score of 0.2. Second, there may clusters which have a dissimilarity score of "1." These are often the clusters which link up at the top of the graph. In terms of membership, these clusters have a have no relationship with one another and form mutually exclusive alliance networks. In Figure 8, there are two alliance networks. Third, following from the previous point, alliances and member states located in different networks will have similarity score of zero while alliances and nations within a given alliance network will have some positive, non-zero similarity score.

Measuring network alliance portfolio similarity

Figure 5 provides a view of the alliance data from the perspective of alliances. Figure 8 does the same but imposes a hierarchical structure based on our measure of similarity. However, in order to measure alliance portfolio similarity between individual nations, we cannot simply create a hierarchical version of Figure 6 since that only involves a change of perspective from alliances to nations. Instead, we need a hierarchical version of Figure 7 which contains information about both alliances and member states.

To do that we need to do some simple matrix manipulation of the data. As mentioned above, the basic alliance data is contained in an incidence matrix with individual alliances as observations (rows) and member nations as variables (columns). What we want is a kind of adjacency matrix in which individual nations are observations (rows) and nations that are in the various alliance portfolios are variables (columns). To get that modified adjacency matrix, I transform the data through the matrix multiplication of the transpose of the incidence matrix multiplied by the incidence matrix itself. This produces a square matrix which has nations as both observations (rows) and variables (columns) but which incorporates the information about alliance relations.[4]

Then, to calculate the measure of network alliance portfolio similarity, I again use hierarchical clustering on this new, transformed matrix. Here, instead of the degree of similarity among alliances as the end product, we now have the degree of similarity among the alliance portfolios of nations. I depict the dendrogram of alliance portfolio similarity for the year 1816 in Figure 9. Here, the vertically oriented two-digit numbers are the last two digits of the four-digit Correlates of War alliance code. The final step is to extract the distances between all pairs of states in the diagram. I do so for the entire Correlates of War alliance dataset for the period from 1816 to 2000.[5]

III NETWORK EFFECTS

Empirical effects of a network perspective

The above discussion may seem plausible, reasonable and maybe even "correct." Nevertheless, the more skeptical may wonder just how much of a difference any of this makes. In truth, this depends on what the alliance data looks like. To that end, I offer four snapshots of the alliance data: 1) the proportion of states that belong to any military alliance; 2) the size of an alliance as measured by the number of member states; 3) the degree of overlapping membership (i.e., the number of alliances to which a nation is a member); and 4) the number of mutually exclusive alliance networks. The first three are positively correlated with any measure of alliance portfolio similarity. The last is negatively correlated.

In Figure 10, one can see the overall degree of alliance participation – the proportion of states that are members of a military alliance is fairly high with approximately 50% of states participating in an alliance in any given year. In Figure 11, one can see that the typical size of an alliance is two states. However, there are some large alliances (approximately 35 members) which drive the average above 3. In

[4] Incidentally, the off-diagonal elements of the resulting matrix is the number alliances of which a given pair of row and column nations are both members. The diagonal elements are the number of alliances of which a specific nation is a member.

[5] The resulting data set is on the order of 1.3 million observations (pair-permutations or directed dyads).

Figure 9 Alliance Portfolio Similarity

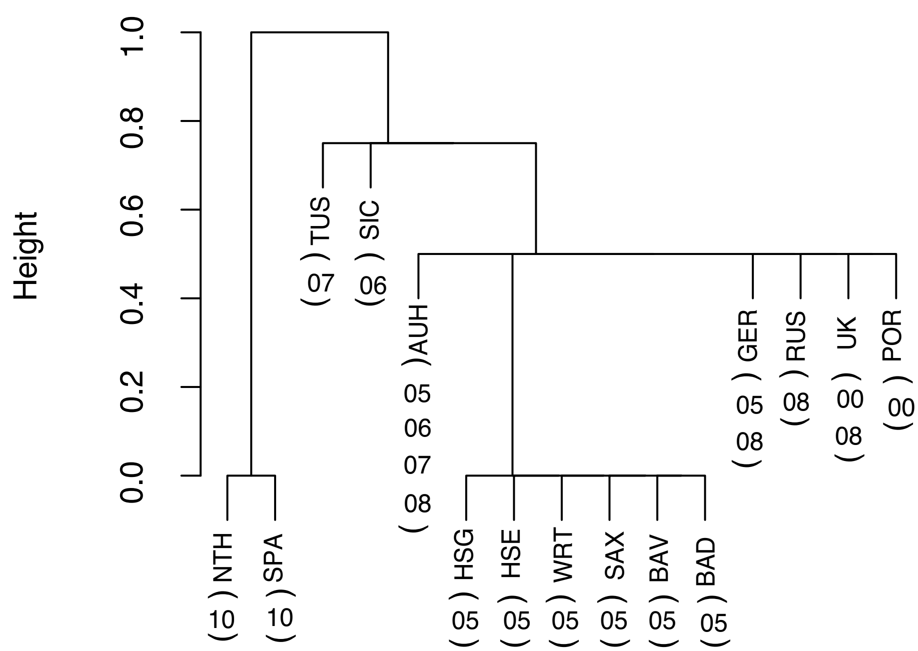

Figure 10 Alliance Participation Rate

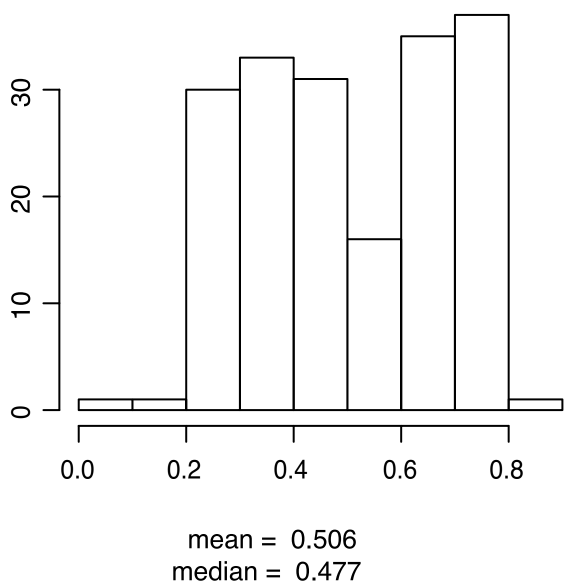

Figure 11 AllianceSize (members)

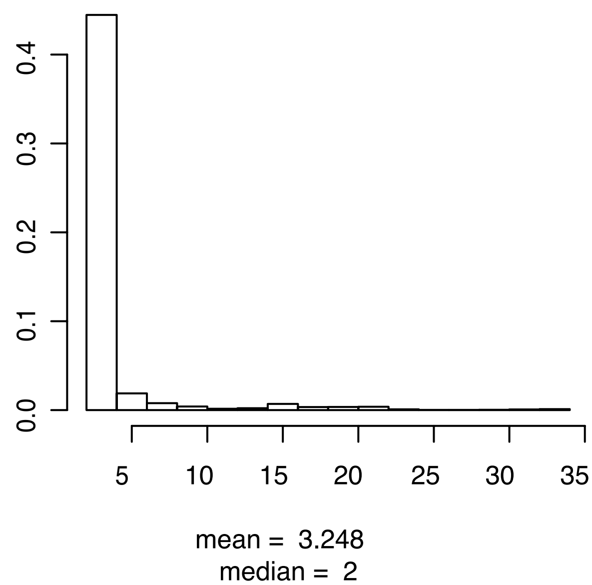

Figure 12 Multiple Alliance Membership

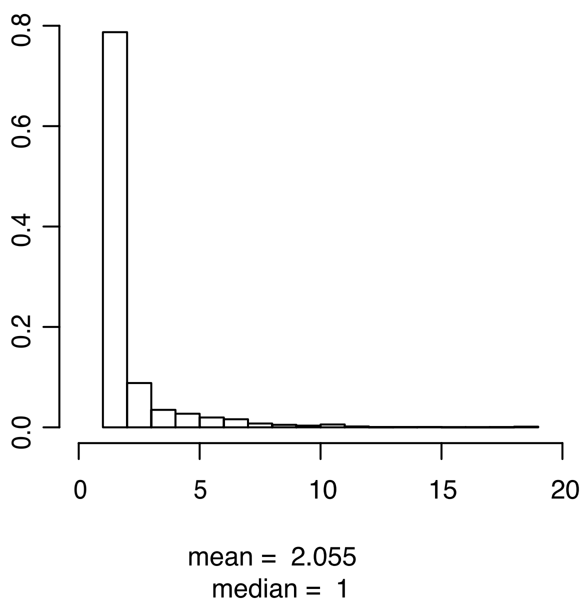

Figure 13 Alliance Networks

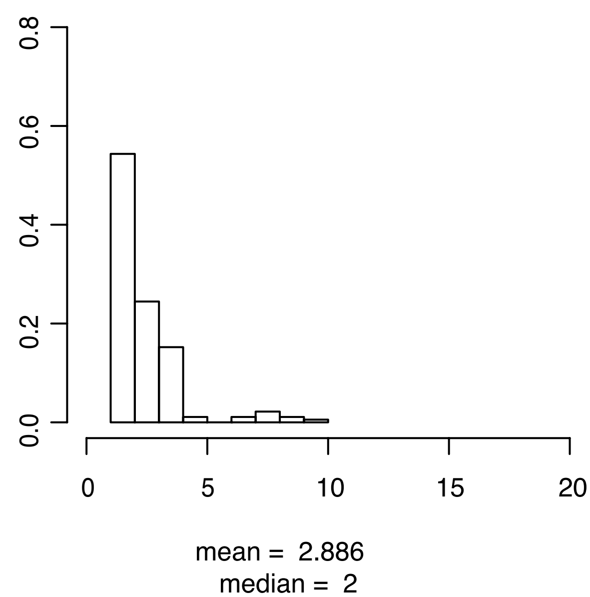

Figure 12, one can see the degree of overlapping membership. While most states are members of a single alliance, there are some large outliers. Finally, in Figure 13, one can see the number of mutually exclusive alliance networks. Since states that are members of different networks have zero similarity, the greater the number of networks, the lower the overall alliance portfolio similarity. The general picture that emerges from these graphs is an overall sparseness in the data which would seem to indicate that overall degree of alliance portfolio similarity should be fairly low.

How well do the three measures, *tau-b*, *S* and my measure of network similarity match this evidence? The answer, illustrated in Figures 14 through 17, is that with the exception of Signorino and Ritter's *S*, the various measures, tau-b, network similarity and the alliance dummy, are at least not inconsistent with the expectation of a low degree of similarity.[6] To see how I reach this conclusion, we need to keep in mind two things. First, with the exception of the dummy variable, the three measures are trying to measure the same thing, alliance portfolio similarity, and use the same data. Second, because *tau-b* and *S* lie on the interval from –1 to 1, and my measure of network similarity lies on the interval from 0 to 1, a direct comparison is not possible. However, the shape of the distributions are comparable. For *tau-b*, a "1" indicates an identical portfolio while a "-1" indicates complete dissimilarity or opposition while "0" indicates independence of rankings. For *S*, "1" indicates an identical portfolio while a "-1" indicates maximum dissimilarity in the sense that portfolios are as far apart as possible. For my measure of network similarity, "1" indicates an identical portfolio while an "0" indicates the absence of any matching members across portfolios.

Based upon these difference, we can conclude that the observed distribution of *S* indicates a high degree of alliance portfolio similarity while *tau-b* and my measure of network similarity indicates a low degree. I would argue that much of *S*'s skewed distribution is the product of the inflation due to the fact that it is a symmetric measure which "overcounts" instances of zero as contributing to alliance portfolio

[6] Data for *tau-b* and *S* are taken from the software management program EUGene. http://www.eugenesoftware.org

Figure 14 Bueno de Mesquita'sTau-b

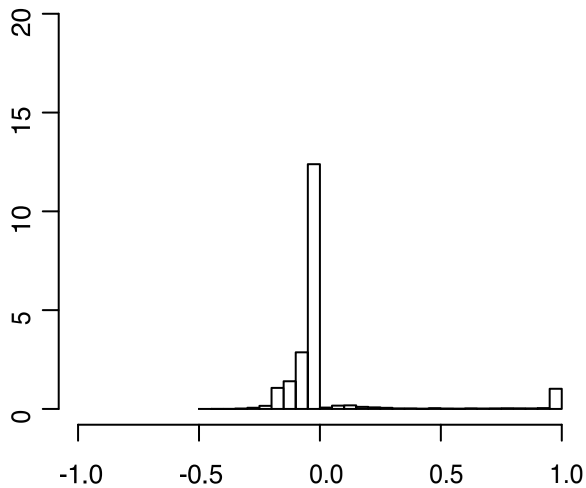

mean = 0.021
median = -0.031

Figure 15 Signorino & Ritter's S

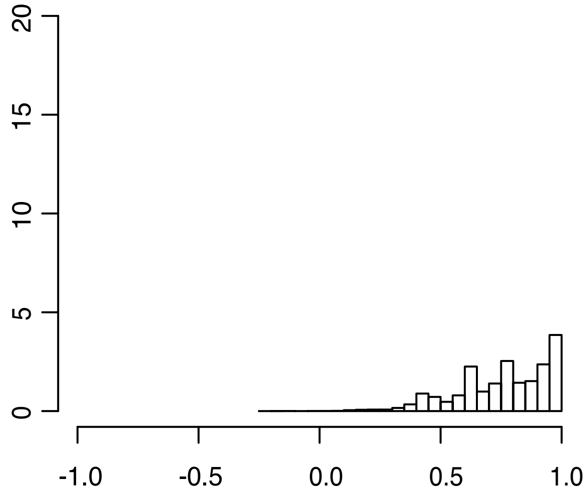

mean = 0.759
median = 0.789

Figure 16 Network Similarity

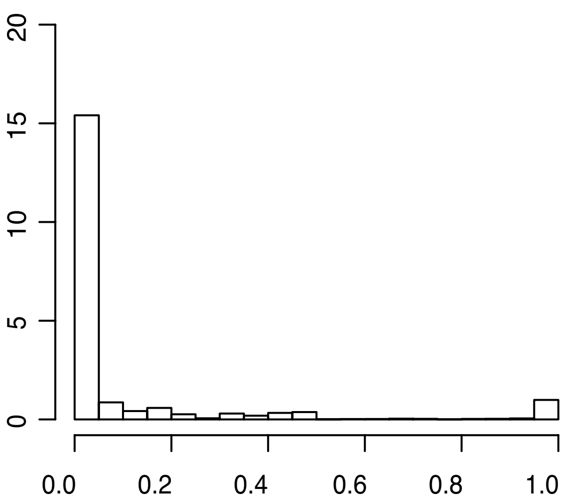

mean = 0.097
median = 0

Figure 17 Alliance Dummy

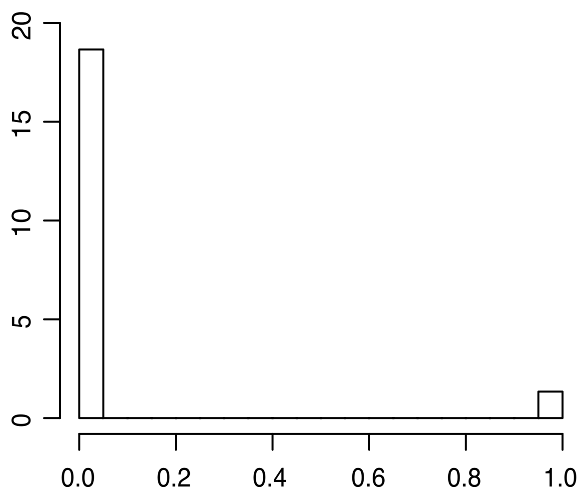

mean = 0.067
median = 0

similarity.[7] The similarity between *tau-b* and my measure of network similarity is also worthy of attention. It is almost as if my measure is a "truncated" version of *tau-b*. Finally, we also observe a similarity in shape between network similarity and the alliance dummy variable. However, none of this mean that *tau-b*, network similarity and the alliance dummy are identical or equivalent.

In order to better assess the relationship between these measures, I plot each against the other in Figures 18 through 23.[8] In each scatterplot, the dotted line links the opposite endpoints of the respective measures' scales. This line represents perfect correspondence between measures. If the two measures were identical, both the data and the solid, fitted OLS line would fall on the dotted line. In Figure 18, we see a some appearance of a structural relationship between *tau-b* and *S*. The commonality, I would surmise, is a product of the fact that both are symmetric measures and that the structure we see is an artifact of the weight of common zeros. In Figure 19, we seem fairly high correspondence between *tau-b* and our measure of network similarity at the high end of their respective scales but a noticeable divergence at the low end. Thus, despite the similarity of the shape of their individual distributions (Figures 14 and 16), the two measures appear not to be identical. The last three plots compare our measure of alliance portfolio similarity against a simple alliance dummy variable which indicates whether or not two states have direct alliance ties. For *tau-b* and *S*, (Figures 21 and 22) there is a noticeable lack of correspondence. The exception is Figure 23 which compares the alliance dummy variable and our measure of network similarity. In contrast to *tau-b* and *S*, since the alliance dummy variable only indicates the common presence of alliance, the difference or the lack of fit is can be attributed to the presence and degree of transitive relations. Generally speaking, the dummy variable underestimates similarity when there is no direct formal alliance between a pair of states and overestimates similarity when there is a direct formal alliance between the pair of states. To get a better sense of how big an effect

[7] Signorino and Ritter argue that to control for this we need to define the domain of "politically relevant" countries. Above in Fn. 3, I argued that this creates more problems than it solves. A robust measure should neither have to rely on the analyst to specify the domain nor be sensitive to such choices.

[8] In contrast to earlier plots, these graphs represent a 5% sample of observations (approximately 65 thousand points). I did so I would not have to plot some 1 million somewhat redundant points. This does not affect any of my inferences, however.

Figure 18 Tau−b v. S

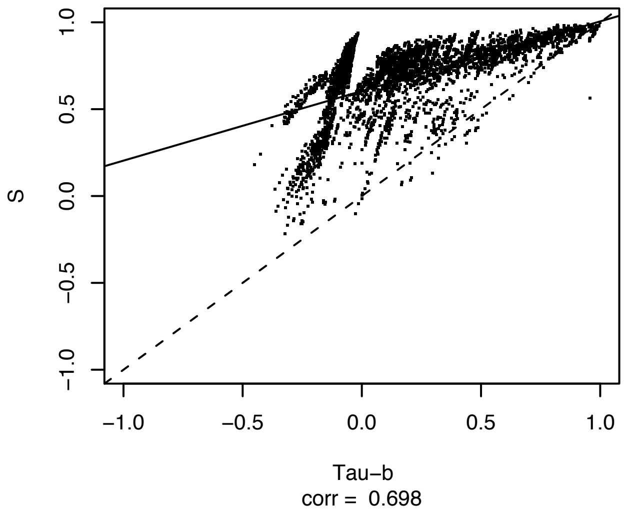

Figure 19 Tau−b v. Network Similarity

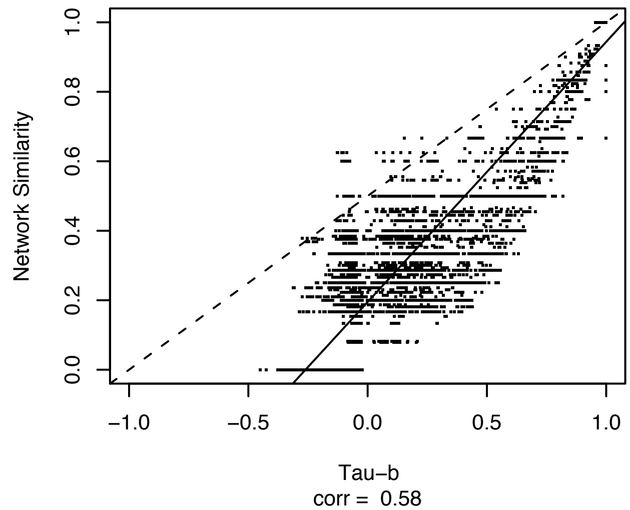

Figure 20 S v. Network Similarity

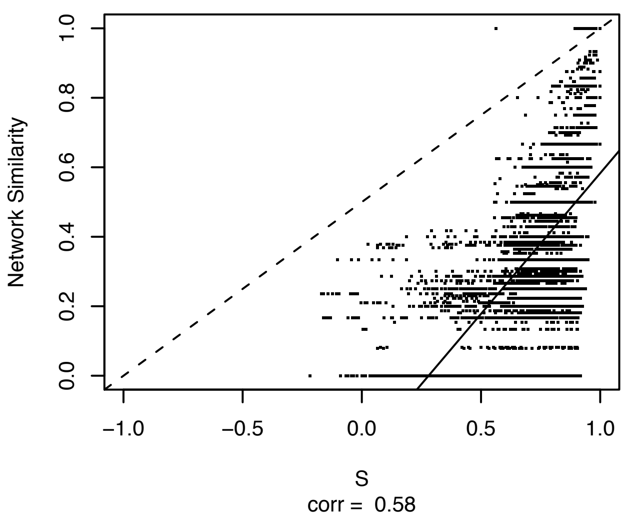

Figure 21 Alliance Dummy v. Tau−B

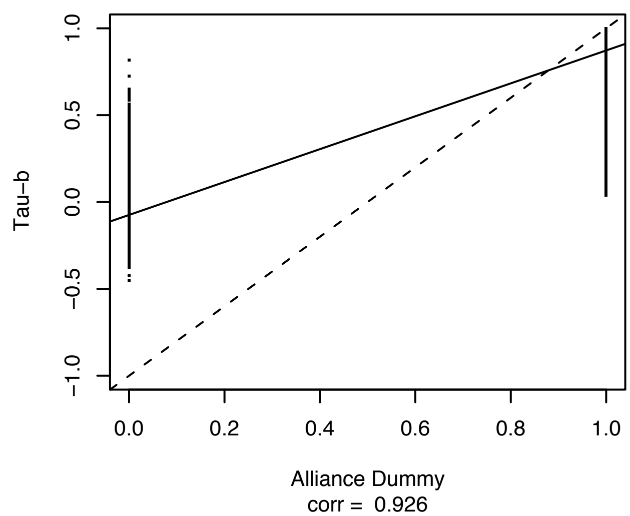

Figure 22 Alliance Dummy v. S

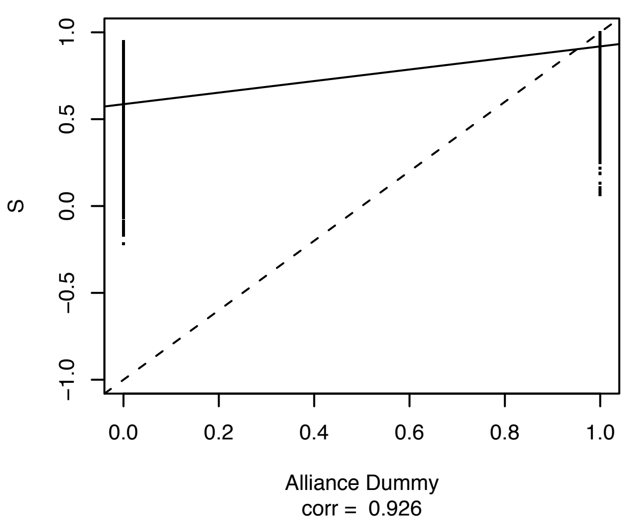

Figure 23 Alliance Dummy v. Network Similarity

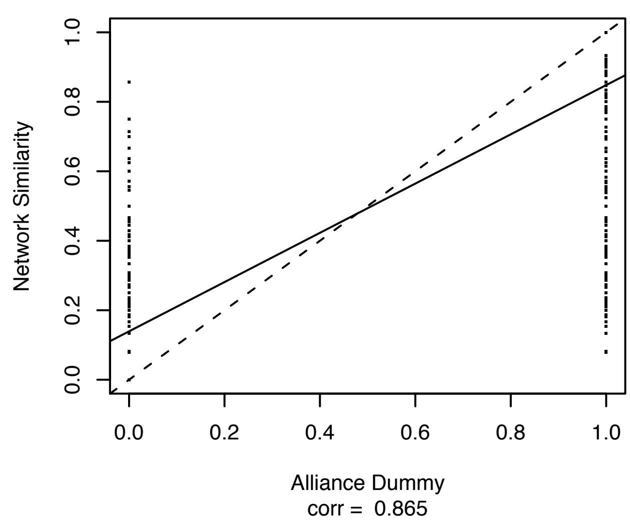

this is, I plot the distribution of network similarity when the alliance dummy equal zero and one as separate histograms in Figures 24 and 25.

Next steps

From these results, one might conclude that to measure the effect of alliances either our measure of network similarity or the simply alliance dummy variable would suffice. However, such a conclusion would be premature. First, we need to assess the inferential and empirical differences of these measures by comparing their performance in a series of standard models explaining phenomena like war or the democratic peace by examining their relative power of out-of-sample prediction. Second, we may need to refine our measure of network similarity. However, this would not involve some exogenous notion of political "relevance" or major power status. Instead, what we may need to do is to normalize the measure according to endogenous differences in the size of alliances (as measured by the number of members). The reason being that larger alliances will simply tend to have lower similarity than smaller alliances. The beneficial side effect is that this method will probably to lead us to the same destination of giving greater weight to important and pivotal alliances and states without the baggage of having to make arbitrary or subjective decisions. Following from the previous point is the third, and perhaps most important point. While our measure of alliance portfolio similarity does much to incorporate the structure of alliance networks into the standard pair-wise view of interaction (i.e., dyad), perhaps it is too optimistic that we could really squeeze in all the relevant information.

To put it another way, perhaps we need to take the network perspective even further. Consider the important role played by Austria-Hungary and Prussia in 1816. In Figure 7, we saw that these nations played an important role as bridges tying together smaller alliances into a larger network structure. The importance of their position at the center or core of that network is not clearly reflected in a measure of alliance portfolio similarity (See their position in Figure 9). In order to really ascertain the effect and importance of alliances in explaining international phenomenon, we may need a better way of

Figure 24 Network Similarity when Alliance Dummy = 0

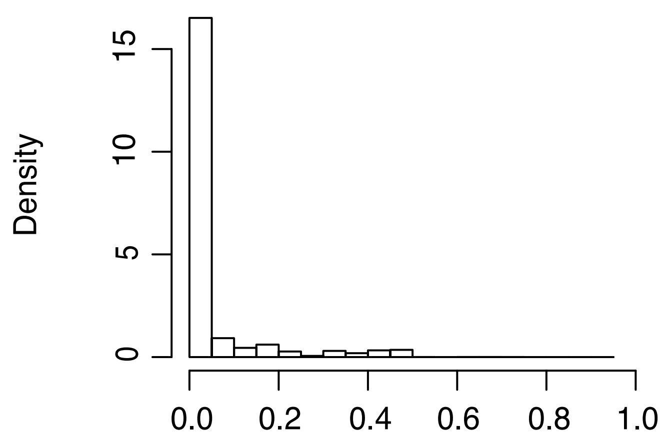

mean = 0.039
median = 0

Figure 25 Network Similarity when Alliance Dummy = 1

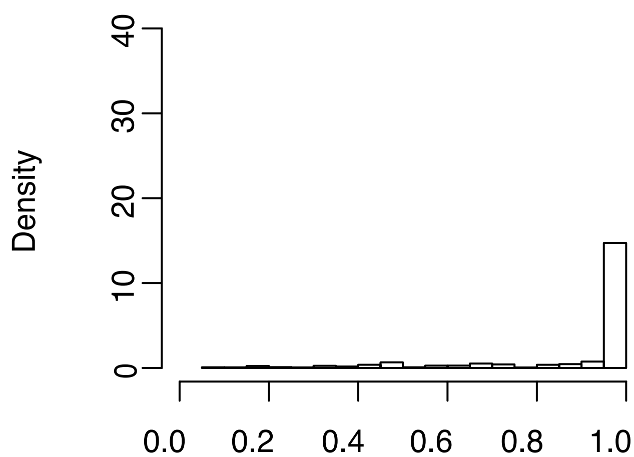

mean = 0.896
median = 1

incorporating the wealth of information reflected in the networks graphs (Figures 5 through 7) than our standard dyadic or systemic perspective currently allow.

IV Conclusions

This paper began with the claim that current images of the effect of interaction in the international relations are either too fine or too coarse. In contrast, I argued that a network perspective captures the happy medium which avoids the sacrifices made by either focusing too much or too little on the structure of relations between actors and events. As a proof of concept, I showed how the most bare-bones notion of a network, one based on the transitive links between alliances can produce a better and more accurate measure of alliance portfolio similarity. My measure of network similarity is better than existing measures because it more closely reflects what the actual alliance data looks like and because it also cleaves more closely to the original conceptualization of the measure as an item-by-item comparison of lists. The fact that the resulting measure also appears to be qualitatively and quantitatively different from existing measures provides some evidence of the value of a network perspective adopted in this paper.

This study of military alliances is but one of many possible applications of the approaches and methods discussed above. More important, perhaps, is the fact that for many applications, this approach is the closest thing we get to a free lunch in terms of new data. All this is really needed is a change in perspective. By the recognizing that there are patterns of relations among observations or actors, we can better leverage existing data.. In short, there is an untapped source of information awaiting our analysis, all that is needed is the curiosity to go out and exploit it.

Bibliography

Michael F. Altfeld and Bruce Bueno de Mesquita. 1979. “Choosing Sides in Wars.” International Studies Quarterly. 23: 87-112.

D. Scott Bennett and Matthew C. Rupert. 2003 “Comparing Measure of Political Similarity: An Empirical Comparison of *S* versus *Tau-b* in the Study of International Conflict.” *Journal of Conflict Resolution*. 47:367-393.

D. Scott Bennett and Allan Stam. 2000. “EUGene: A Conceptual Manual.” International Interactions 26:179-204. http://www.eugenesoftware.org

Bruce Bueno de Mesquita. 1975. “Measuring Systemic Polarity.” *Journal of Conflict Resolution*. 23: 187-216.

Ronald S. Burt. 1992. *Structural Holes: The social structure of competition*. Cambridge: Harvard University Press.

James D. Fearon. 1994. “Domestic Political Audiences and the Escalation of International Disputes.” *American Political Science Review*. 88:577-592.

Leonard Kaufman and Peter J. Rousseeuw. 1990. *Finding Groups in Data: An Introduction to Cluster Analysis*. New York: John Wiley & Sons, Inc.

M.E.J. Newman. 2003. “The Structure and Function of Complex Networks.” *Society for Industrial and Applied Mathematics Review*. 2003. 45:167–256.

Keith T. Poole and Howard Rosenthal. 1991. “Patterns of Congressional Voting.” *American Journal of Political Science*. 35:228-78.

Kenneth A. Schultz. 2001. *Democracy and Coercive Diplomacy*. New York: Cambridge University Press.

Curtis S. Signorino and Jeffery M. Ritter. 1999. “Tau-b or Not Tau-b: Measuring the Similarity of Foreign Policy Positions.” *International Studies Quarterly*. 43:115-144.

Lawrence Sirovich. 2003. “A pattern analysis of the second Rehnquist U.S. Supreme Court.” *Proceedings of the National Academy of Sciences*. 100: 7432-7437.

Kenneth N. Waltz. 1959. *Man, the State, and War*. New York: Columbia University Press.

Kenneth N. Waltz. 1979. *Theory of International Politics*. New York: Random House.

Stanley Wasserman and Katherine Faust. 1994. *Social Network Analysis: Methods and Applications*. New York: Cambridge University Press.