Characterizing Social Fragmentation over Social Networks: An Axiomatic Approach

Neelanjan Sircar

September 20, 2014

Abstract

This paper develops two axiom-based measures to characterize social fragmentation over social networks: generalized fractionalization and proclivity. Many traditional measures of fragmentation, like ethnolinguistic fractionalization (ELF), are highly sensitive to researchers' definitions of groups and social categorizations and thus highly susceptible to researcher biases. The measures discussed in this paper do not require researchers to define groups a priori if a social network can be observed or perceived between members of the population. In particular, the paper proposes a series of intuitive axioms that uniquely characterize the suggested measures, thereby providing an intuitive meaning to fairly objective measures over a social network. Furthermore, it is shown that these measures can be used to endogenously define social cleavage structures for any social network. In short, this paper develops intuitive axiomatically-characterized measures of social fragmentation which can be derived directly from the structure of a social network without relying upon researcher biases about social categorization.

1 Introduction

This paper develops two axiom-based measures to characterize social fragmentation over social networks: generalized fractionalization and proclivity. Many traditional measures of fragmentation, like ethno-linguistic fractionalization (ELF), are highly sensitive to researchers' definitions of groups and social categorizations and thus highly susceptible to researcher biases. The measures discussed in this paper do not require researchers to define groups a priori if a social network can be observed or perceived between members of the population. In particular, the paper proposes a series of intuitive axioms that uniquely characterize the suggested measures, thereby providing an intuitive meaning to fairly objective measures over a social network. Furthermore, it is shown that these measures can be used to endogenously define social cleavage structures for any social network. In short, this paper develops intuitive axiomatically-characterized measures of social fragmentation which can be derived directly from the structure of a social network without relying upon researcher biases about social categorization.

Section 2 introduces the reader to non-directed networks and notions of homophily, and provides motivation for the proposed definition of social fragmentation. Section 3 develops a uniquely characterized generalized fractionalization measure from three axioms, and section 4 uses similar axiomatic and statistical techniques to measure identity-based clustering in a network and uses these measures to construct a partition of the space. Section 5 concludes the paper.

2 Preliminaries

Traditional methods in assessing identity-based social fragmentation try to characterize the salient identity groups in social fragmentation. The most well-known of these measures is ethno-linguistic fractionalization (ELF). ELF (originally designed by *Atlas Narodov Mira* (1964)) has been subject to much controversy and debate. In fact, there have been recent attempts to create more accurate ELF measures (Posner, 2004; Laitin and Posner, 2001). The classic ELF measure is a Herfindahl index using relative sizes of each group population as inputs in the index.[1] Furthermore, as Fearon (2003) notes, two countries, one with two groups of equal size and one with three groups with shares $\frac{2}{3}$, $\frac{1}{6}$, and $\frac{1}{6}$, would both have an ELF of 0.5.

Measurement of social polarization, i.e., measures of variance and clustering in social connectedness, is less common. One notable exception is the work of Esteban and Ray (1994), who construct a general class of polarization measures from a set of core axioms (much like this paper). In their construction, individuals are endowed with an underlying attribute (e.g., income) upon which social connected/alienation is characterized. This is a measure that is well suited to measuring certain "types" of polarization, i.e., polarization due to income classes, but it maybe difficult to implement when the type of polarization results from an complex combination of multiple attributes since it puts the onus on the researcher to find such a measure.

The bigger problem with these measures is that it requires the researcher to determine "relevant" groups for the measure.[2] For instance, the original characterization of ELF did not view Hutus and Tutsis as separate ethnic groups because they spoke the same language. These measures stand in direct contrast to those who emphasize that people may have multiple identities at once. Sen (2006), in response to anti-Muslim tensions, notes, “The increasing tendency to overlook the many identities that any human being has and to try to classify individuals according to a single allegedly pre-eminent religious identity is an intellectual confusion that can animate dangerous divisiveness." Similarly, Linz and Stepan (1996) argue

[1] The Herfindahl index formula is: $ELF = 1 - \sum_{i=1}^{n} s_i^2$ where s_i is the share for each group $i \in \{1, 2, ..., n\}$.

[2] Daniel Posner's PREG index is one attempt to rigorously define relevant groups by searching newspapers and country-specific data.

that individuals may have multiple and complementary identities, and thus characterization by unique identities tend to overstate the level of fractionalization in society.

At the same time, if individuals hold on to multiple identities at once, we are left with question of when and why certain identities are actionable for the purposes of conflict, i.e. why do people of certain sets of identities fail to build social ties with each other? The approach undertaken in this paper is shy away from a classification of multiple identities and when they may become relevant. Rather, this paper uses underlying or perceived patterns of interaction in a social network, which are necessarily a function of multiple identities and personal attributes, to characterize fragmentation.

2.1 Observing Social Fragmentation

Lack of social ties may occur for two reasons: 1) preferences of individuals, and 2) lack of opportunity to interact. When we speak of social *fragmentation*, we are primarily concerned with the former reason, and when we speak of social *segregation* we usually mean the latter reason. Of course, fragmentation can lead to segregation and vice versa, but it is useful to keep the two concepts analytically distinct. An example should clarify the distinction. Consider the bordering neighborhoods of the Upper West Side, with a large white population, and Harlem, with a large black population, in New York City, which tell us that the city is segregated. Now, as a thought experiment, imagine that populations are mixed so that over the course of the day, any person from either of the two neighborhoods has an opportunity to interact with everyone else in the Upper West Side and Harlem. Presumably, some members of the white and black populations would form social ties with each other when they did not have the opportunity to do so before. On the other hand, there will be those who refuse to form social ties across race, even with guaranteed interaction. We can now put the populations in the following equation:

$$\text{(\# No Social Ties Before Mixing)} =$$

$$\text{(\# No Social Ties After Mixing)} + \text{(\# Social Ties After Mixing, but No Ties Before Mixing)}$$

The left side of the equation refers to the estimate of the lack of social ties under segregation, and the first term on the right refers to what I shall call social fragmentation, *the lack of social ties between individuals when given the opportunity to interact*. We can see that observing social structure under a condition of segregation overestimates fragmentation.[3] Interaction is the crucial element here. In particular, interaction allows individuals to individuals to assess multiple identities at once and make choices about cooperation based upon these assessments. As Goffman (1983) writes,

> Once individuals–for whatever reason–come into another's immediate presence, a fundamental condition of social life becomes enormously pronounced, namely, its promissory, evidential character. It is not only that our appearance and manner provide evidence of our statuses and relationship. It is also that the line of our visual regard, the intensity of our involvement, and the shape of our initial actions, allow others to glean our immediate intent and purpose, and all this whether or not we are engaged in talk with them at the time (p. 3)

In this view, the best approach to measuring social fragmentation is to directly observe interaction between individuals. Of course, this sort of observation is not always feasible; however, it does provide us with a benchmark or ideal for the measurement of social fragmentation. Nonetheless, the measures described in this paper are well-defined even without the opportunity/preference distinction, but the interpretation of such social fragmentation differs based on the relative opportunity of individual to interact with each other.

[3]The assumption here is that people do not break social ties after mixing. In a situation where this occurs, there may be fewer social ties under mixing than under segregation.

2.2 Games of Interaction and Coordination

In the study of sociology, economics and politics, we are often interested in how our objects of study behave when they are forced to interact with each other. The logic of these situations are different from the sort of community-mapping exercise that is common in sociology. Inferences from community-mapping are interested in the proximity of individuals with each other. Taking this as a starting point, the goal is to understand the probability of interaction between two individuals or the flow of information in a network. In the setup for this paper, we assume that every individual has an opportunity to "meet" every other individual in the study.

A number of situations can be modeled in this way. For instance, in bilateral trade or treaty negotiations, both countries must agree. However, theoretically, no country is a priori restricted from interacting with another country. Another situation would be the co-authorship of bills in a legislature. Once again, legislators are perfectly free to interact with each other, but who do they choose to coordinate with on bills?

A second class of situations of interest are when individuals who do not normally interact with each other are forced to do so. In a game of coordination, we are interested in isolating lack of interaction due to stigma, as opposed to lack of opportunity. Individuals use signals available to them in tandem with preconceived notions to decide with whom to coordinate. For instance, we could be interested in investigating how students from disparate "ethnic groups" might interact with each other in a newly integrated school. Or, we might be interested in testing whether such groups would be amenable to supporting the same political party or candidate. Each of these hypothetical scenarios would shed light upon the level of social fragmentation in the sample.

2.3 Understanding Non-Directed Graphs

The theory in this paper is based on an analysis of non-directed graphs, or social networks. Non-directed graphs model pairwise interaction. We may view ties or links in a network as coordination between two individuals, as in friendship or a shared task. Typically, we view this coordination as stemming from some social “similarity" between the individuals, or *homophily*. The graph is considered “non-directed" because both individuals must agree in order for cooperation to take place. In other words, links in a non-directed graph indicate reciprocal behavior and agreement.

Many of the situations we have discussed above can be modeled with non-directed networks. Consider the example of friendship networks in schools. A "link" is formed between the two individuals in a dyad if each individual views the other as a friend, i.e. beliefs of friendship are reciprocal. Although each person may use some combination of race, gender, or attractiveness to determine friends, the link is solely a function of interaction between the individual in the dyad. Thus, the researcher places no restriction in identity categorization upon the individual in order to determine cooperation or connection. Figure 1 shows a set of possible graphs formed by this routine.

Graph A is what we might think of as a divided system, where there are two components, each of which is a complete subgraph (i.e. each individual has formed a link with each other individual in the component). The level of connectedness measures in some way the density of the group structure. So, for instance, graph A denotes two very dense groups that do not interact with each other. Graph B shows that a connected graph need not change the structure too much, but now we don’t have two distinct groups. Imagine that the left group is the "black" group and right group is the "white" group. The link between the two groups suggests that there is a person who is mixed in a way that suggests that she is accepted by both black and white groups and she views herself as both black and white. However we still basically see two groups here. Since the structure is quite similar, we want fractionalization in A and B to be similar, a topic discussed in the following section. Graph C shows a situation where one individual is isolated from to tight-knit groups, and graph D shows a situation where there aren’t two easily discernible groups

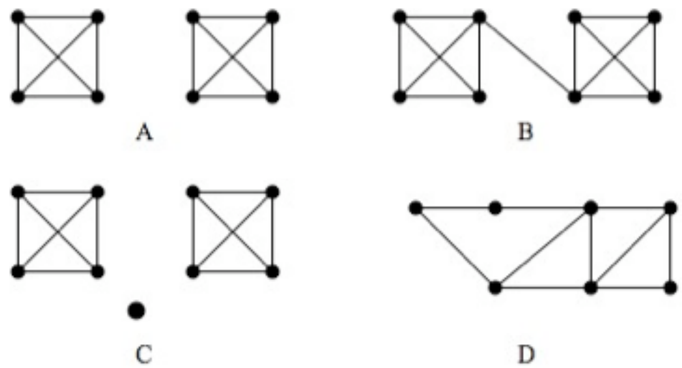

Figure 1: A Few Possible Networks

and the network is not as dense in clusters as graph A. While networks are an extremely flexible way of representing social structure, they are also extremely difficult to summarize statistically.[4] Our task in the following sections is made easier by the fact that we are working in a context where each individual is allowed to interact with every other individual.

3 Fractionalization

In common parlance, when we speak of Òsocial fragmentation,Ó we are actually conflating two very different phenomena. Let's consider graphs A and D once again from figure 1. Which network is more fragmented? On one hand, the groups in graph A are much more dense, so we may be inclined to argue that A is less fragmented than D. On the other hand, A exhibits two discernible groups that do not interact with each other, whereas this sort of separation is not present in D. Social fragmentation can be decomposed along two dimensions: fractionalization and variance in connectedness. In this section, we discuss fractionalization, which we intuitively define as the average level of disconnectedness in the network. The issue of variance in connectedness is addressed in the following section.

3.1 Defining Fractionalization

Throughout the paper, we use $G = \{N, V\}$ to denote the social network where$N$ to denote the set of vertices or nodes in the networks V denotes the set of ordered pairs in node set that have links between them. In ELF, the the proportion of group j, s_j, in the whole population is used in the formula. Here, we define a proportion, *connectedness*, for each individual.

Definition 3.1. *Let δ_i denote the degree (number of links) for person i. The personal connectedness of player i, p_i is just $\frac{\delta_i}{n-1}$, where $|N| = n$.*

In turn, the generalized ethno-linguistic fractionalization (Z) is just a function of p_i:

Definition 3.2. *The generalized ethno-linguistic fractionalization, Z is given by the following formula:*

$$Z = 1 - \frac{1}{n}\sum_{i=1}^{n} p_i = 1 - \frac{1}{n(n-1)}\sum_{i=1}^{n} \delta_i$$

[4]The standard text on different statistical measures to characterize networks is Wasserman and Faust (1994).

One natural interpretation here is that Z represent the fraction of dyads that did not form a link between them. In other words, Z represents the probability that two randomly chosen individual do not have a link between them. The index, Z, ranges from 0 to 1, and clearly the maximum value of the index goes to 1 as the sample size grows. Furthermore, Z is monotonically decreasing in the total number of links formed in the network. The expression is remarkably simple, as it is a function of the average of connectedness measures in population. Note that the definition is of connectedness here is quite different from the definition in Krackhardt (1994), which is more concerned with the reachability from one individual to another, i.e. there exists a path that connects two nodes.

Figure 2: Various Social Structures

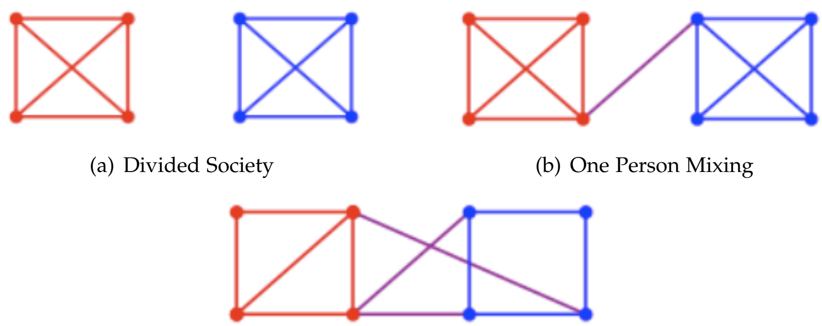

(a) Divided Society

(b) One Person Mixing

(c) Many People Mixing

The graphs above depict two groups, the"red" and "blue" groups. When linking within group, the link is the color of the group. When linking across groups, the link is purple. Even though (a) and (c) have very different structures, they admit the same generalized fractionalization score, $\frac{4}{7}$ or 0.571 in the frequentist measure, or $\frac{17}{30}$ or 0.567 in the Bayesian measure (see section 3.3 for discussion).

In figure 2, (a) and (c) have identical values of Z , 0.57. Since the fractionalization score is simply conceived as average disconnectedness, it does not account for any clustering behavior. Although, Z is not a function of experimenter coding like the ELF, it is still subject to the same non-uniqueness problem of the classic measure. In fact, there is an intimate connection between ELF and Z. Consider a network divided into m components[5], then, as the following theorem shows, ELF and Z correspond under special circumstances:

Theorem 3.3. *Let G_n denote a graph of sample size n, and let Z_n correspond to the value of Z on G_n. Assume that each G_n contains m (fixed) components, each of which is a complete subgraph. Then,*

$$\lim_{n\to\infty} Z_n = ELF$$

Proof: Assume G_n contains m components, each of which is complete, and s_j corresponds to the fraction of the population in component j, then:

$$\lim_{n\to\infty} Z_n = \lim_{n\to\infty} 1 - \frac{1}{n(n-1)}\sum_{i=1}^{n}\delta_i = \lim_{n\to\infty} 1 - \frac{n}{n-1}\sum_{j=1}^{m} s_j^2 - \frac{1}{n-1} = 1 - \sum_{j=1}^{m} s_j^2 = ELF$$

From this perspective, the classical *ELF* makes very strong assumptions about the level of connectedness in a network. It implies that members of one group are connected with every other member of the group and no one else.[6] Thus, we can surmise that the *ELF* assumes within-group connectedness is quite dense

[5]A component is defined as the maximal subgraph where each pair of nodes can be connected by a path. In other words, two individuals are in different component if there is no set of links connecting in the graph connecting the individuals to each other.

[6]Of course, there may be knife-edge results where members of the group are not completely connected to each other, but are connected outside the group in such a proportion to make the total number of links correspond approximately to what would have occurred in the assumptions of theorem 3.2.

and across group connectedness is very sparse in large samples. However, the direction of the bias is ambiguous. For instance, it may be that within-group connectedness is more sparse than assumed under classical ELF. In this case, the ELF underestimates the level of fractionalization in the population. On the other hand, there is certainly some connectedness across groups, and ELF overestimates fractionalization in that situation. Another common interpretation of the ELF is that it measures the probability that two people chosen at random will belong to different groups, whereas Z measure the probability that two people chosen at random do not have a link.

3.2 An Axiomatic Approach to Fractionalization

In order to further understand the measure proposed above, this subsection attempts to construct the measure from an intuitive set of axioms. These axioms constitute a sufficient set of conditions to characterize a network using Z. The validity of Z is a function of the intuitive appeal of the axioms presented below. In this subsection, we will present each axiom followed by a brief discussion, concluding with the main theorem of the subsection.

In order to orient the discussion, we define a process of graph formation. Assume that the set of individuals, N, has cardinality n. N induces a dyad set, D, of cardinality $\frac{n(n-1)}{2}$ with some ordering. A dyad (pair of individuals) in D is chosen randomly, one at a time without replacement, until all $\frac{n(n-1)}{2}$ dyads have been selected. The result of the process is a $\frac{n(n-1)}{2}$-tuple denoted by D^*. For In each period, the dyad selected is given the opportunity to form a link. Our goal is to determine a function $f : P(D^*) \rightarrow [0, 1]$ that assigns the space of permutations of D^*, $P(D^*)$ to a number between 0 and 1, defining the fractionalization index. The following definition is helpful for characterizing the axioms:

Definition 3.4. *Let $A \subseteq D$ be the subset of dyads that have formed links. We define $D^*(A)$ to be a process where the members of A have formed links.*

The first of our axioms is monotonicity, which requires that forming a new link necessarily decreases the fractionalization. The second axiom, anonymity, requires that each dyad has the same effect on the index. The third axiom, order invariance, requires that dyads have the same effect on the index irrespective of when they are chosen. The axioms are stated formally below.

Monotonicity. *Let $A, B \subseteq D$ be the set of dyads that have a link between them in two different processes, $D^*(A)$ and $D^*(B)$. If $A \subsetneq B$, then $f(D^*(A)) > f(D^*(B))$.*

Anonymity. *The fractionalization function, f, is not dyad-specific. Formally, we say $f\left((\Pi D)^*\right) = f(D^*)$ where ΠD is a permutation of D.*

Order Invariance. *Dyads have same impact on the index, irrespective of when they are chosen. Formally, we say $f(\Pi D^*) = f(D^*)$ where ΠD^* is a permutation of D^*.*

The appeal of monotonicity is fairly obvious. As more people form links, the fractionalization in society decreases. Notice, however, that the condition above is weaker than stating a process with more links has lower fractionalization. We require that the sets of dyads forming links in two processes be nested in order to satisfy the monotonicity condition. Intuitively, we may take a process and add a link between a dyad that failed to form one, which creates a new process that is strictly less fragmented than the old one. The anonymity condition is an egalitarian condition that allows us to calculate the same index irrespective of *who* is forming a link. The order invariance condition can be thought of as an independence condition. Link formation between two individuals is not a function of the actions of other dyads and thus the contribution of a link between the elements of a dyad in the fractionalization index is independent of the actions of the other links. These three axioms allow us to derive the main result of this subsection:

Theorem 3.5. *The fractionalization index, Z, is the unique function $f : P(D^*) \to [\alpha, \beta]$ that satisfies monotonicity, anonymity, and order invariance with a range of $[0,1]$.*

Proof: Consider a process D^* and some dyad $k \in D$. The marginal effect on f of forming a link in the process is defined by:

$$f\left(D^*(A \cup \{k\})\right) - f\left(D^*(A)\right) = c_k$$

By order invariance, the marginal effect of forming a link for k is a constant c_k irrespective of when it is chosen and thus dyad k has a marginal effect of c_k in every process. By anonymity, however, $c_k = \alpha$ for all k, i.e. each dyad has the same contribution. It follows that f is linear in the number of links, so:

$$f(D^*) = \beta - \alpha \times \sum_{k \in D} d_k \quad \text{where } d_k = \begin{cases} 0 & \text{if no link is formed} \\ 1 & \text{if a link is formed} \end{cases}$$

Monotonicity implies that $\alpha > 0$ and a range of $[0,1]$ implies that $\beta = 1$ and $\alpha = \frac{2}{n(n-1)}$. Now, we use the fact that:

$$2 \times \sum_{k \in D} d_k = \sum_{i \in N} \delta_i$$

Substituting in directly gives us:

$$f = 1 - \frac{1}{n(n-1)} \sum_{i=1}^{n} \delta_i = Z \quad \square$$

3.3 Stochastic Approaches to Fractionalization

In the previous subsection, the axiomatic approach constructed the fractionalization index as a deterministic quantity. However, we can also treat the index as a stochastic quantity based on the probability definition, the probability that two people chosen at random do not have a link between them. Accordingly, we can give both a frequentist and Bayesian interpretation of the measure.

Why might we select one over the other? Apart from philosophical concerns about one interpretation or the other, the main difference between the two interpretations is about how to handle sample size. The frequentist approach simply counts the number of links and divides by the number of possible links to get probability of link formation. Thus, the frequentist measure is precisely Z as we have already defined it.

One Bayesian critique of this approach is that it give non-credible results in smaller sample sizes. Imagine that we only sample two people from a population and they form a link. Do we really believe that the probability of not forming a link in this population is 0? The Bayesian measure described here starts off with a non-informative prior, which says that prior belief about probability of link formation (and, thus, non-formation) is uniformly distributed between 0 and 1, which intuitively pulls probability of link formation (or non-formation), also known as the posterior, towards $\frac{1}{2}$. This allows us to give the Bayesian version[7] of Z below:

[7]For more discussion of the Bayesian estimator(s), please see Gelman, Carlin, Stern and Rubin (2003)

Definition 3.6. *Let the prior belief about the probability of link formation in a population be uniformly distributed on* $[0, 1]$*. Then, the Bayesian version of the fractionalization index is given by:*

$$Z^B = 1 - \frac{(\sum_{i=1}^{n} \delta_i) + 2}{n(n-1) + 4}$$

The standard errors of the two types of estimates are fairly straightforward to calculate. In the frequentist case, we have the probability of a link $p = \frac{\sum_{i=1}^{n} \delta_i}{n(n-1)}$. In the frequentist case, the data are generated from a binomial distribution with paramater p, and thus standard error of the frequentist estimate of Z is just:

$$se(Z) = \frac{\sqrt{p(1-p)}}{\sqrt{\frac{n(n-1)}{2}}} = \frac{2 \times \sqrt{n(n-1) \sum\limits_{i=1}^{n} \delta_i - \sum\limits_{i=1}^{n} \delta_i}}{n^{\frac{3}{2}}(n-1)^{\frac{3}{2}}}$$

. It can be shown that probability of link formation in the Bayesian case, where the probability is just $p^B = \frac{(\sum_{i=1}^{n} \delta_i)+2}{n(n-1)+4}$, follows a beta distribution with parameters $\frac{\sum_{i=1}^{n} \delta_i}{2} + 1$ and $\frac{n(n-1)}{2} - \frac{\sum_{i=1}^{n} \delta_i}{2} + 1$. It follows that estimated standard error of Bayesian case, Z_B, is just:

$$se(Z^B) = \frac{\left(\frac{\sum_{i=1}^{n} \delta_i}{2} + 1\right)\left(\frac{n(n-1)}{2} - \frac{\sum_{i=1}^{n} \delta_i}{2} + 1\right)}{\left(\frac{n(n-1)}{2} + 2\right)^2 \left(\frac{n(n-1)}{2} + 3\right)} = \frac{2\left(\sum\limits_{i=1}^{n} \delta_i + 2\right)\left(n(n-1) - \sum\limits_{i=1}^{n} \delta_i + 2\right)}{(n(n-1)+4)^2 (n(n-1)+6)}$$

In this paper, I do not advocate one measure over the other. Rather, because they are both very easy to calculate, it may be beneficial to calculate both measures and be concerned if the two measures are very far off. In the case where the values of the indices are far off, we would suspect that the sample size is too small, and the frequentist measure should be interpreted with caution.

4 Proclivity

In the previous section, we saw that the fractionalization index alone does not adequately capture major changes in social structure. The approach in this subsection is to derive a difference of probabilities measure which captures group-level clustering. In particular, we may have groupings in mind according to some identity-based measure like race or gender, and we would like to test whether there is clustering interaction across these identities. In this section, I define and discuss the properties of a *proclivity* measure, which is based on the Z measure above. Again, we use an axiomatic approach to derive the measure as a deterministic quantity, and we discuss its stochastic application thereafter.

4.1 Defining and Axiomatizing Proclivity

As we will see, the definitions described in this subsection are straightforward applications of the routine described in the previous section on fractionalization. The statistical estimation of these values, however, can be complicated (which is discussed in the next subsection).

Let the fractionalization of group X restricted to group Y be defined in the following way:

Definition 4.1. *The fractionalization of group* X *restricted to group* Y *is defined as:*

$$Z_{X,Y} = 1 - \frac{1}{n_Y(n_Y - 1)} \sum_{i \in X} \delta_i^Y$$

where $n_Y = |Y|$ and δ_i^Y is the degree of person i restricted to when restricted to members of group Y.

In turn, the proclivity of a member of group X towards a member of group Y (as opposed to a member of the complement of Y, $\neg Y$) is defined by a difference of these restricted fragmentation indices:

Definition 4.2. *The proclivity of group X towards group Y is defined as:*

$$Q_{X,Y} = Z_{X,Y} - Z_{X,\neg Y} = \frac{1}{n_Y(n_Y - 1)} \sum_{i \in X} \delta_i^Y - \frac{1}{n_{\neg Y}(n_{\neg Y} - 1)} \sum_{i \in X} \delta_i^{\neg Y}$$

where, $\neg Y$ denotes the complement of Y, $n_Y = |Y|$, $n_{\neg Y} = |\neg Y|$, and δ_i^Y, $\delta_i^{\neg Y}$ is the degree of person i restricted to when restricted to members of group $Y(\neg Y)$.

Intuitively, this measure is just the probability of a member of X forming a link with someone in Y, subtracting the probability of a member of X forming a link with someone outside of Y, which ranges between -1 and 1. An interpretation of Q would be the “correlation" in link formation between members of groups X and Y. One way to conceive of the measure is to think of “latin square design" experiment. In figure 3, we see a diagram of this scenario. Intuitively, we can partition the dyad set into X and $\neg X$, as well as Y and $\neg Y$. Then, we can further conceive of the probability of links forming in each of the combinations ($< X, Y >$, $< X, \neg Y >$, $< \neg X, Y >$, and $< \neg X, \neg Y >$). The proclivity measure is just the difference of the top two cells (those links in the complement of $\neg X$).

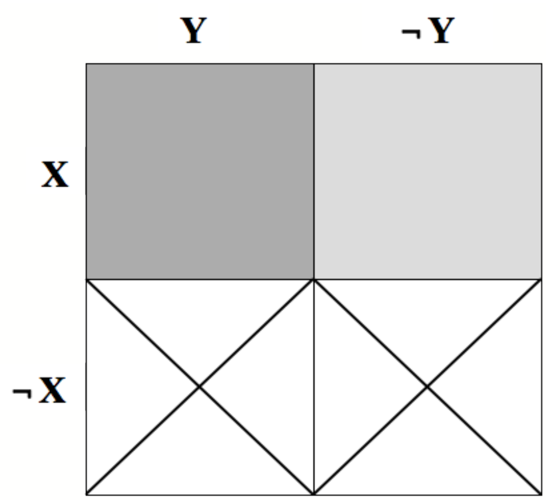

Figure 3: A Latin Square Conceptualization of $Q_{X,Y}$

The latin square image above depicts the set of nodes divided in two ways, by X and its complement, and by Y and its complement. The darker shade implies a higher probability of forming a link, and $Q_{X,Y}$ is just the difference of the two shaded regions.

We may look at the problem cases from before, where very different structures admit the same fractionalization score. In order to test "in-group" clustering, we would measure $Q_{X,X}$. As the next theorem shows, $Q_{X,X}$ has a very intuitive interpretation.

Theorem 4.3. *If $Q_{X,X} = 1$, then X forms a complete component, i.e. each member of X has a link to every other member of X and no one else.*

The proof is trivial, and is thus omitted. However, $Q_{X,X}$ can be thought of as a measure of how close a comes to the extreme case of forming a “clique" (complete subgraph) without connecting to anyone else. Thus, the proclivity measure combines a notion of clustering with a notion of isolation, and groups with high values of Q are dense and isolated.

Figure 4: Various Social Structures

(a) Divided Society (b) Many People Mixing

The graphs above depict two groups, the"red" and "blue" groups. When linking within group, the link is the color of the group. When linking across groups, the link is purple. Even though (a) and (b) have very different structures, they admit the same generalized fractionalization score, $\frac{4}{7}$ or 0.571 in the frequentist measure, or $\frac{17}{30}$ or 0.567 in the Bayesian measure. However, they admit very different proclivity scores, where $Q_{X,X}$ is 1 (frequentist) or 0.9997 (Bayesian) for (a), and where $Q_{X,X}$ is 0.4791 (frequentist) or 0.4783 (Bayesian) for (b). See section 4.2 for a discussion of the Bayesian measure.

The axiomatic construction of proclivity is similar to the one given in the last section, and one of the axioms is identical (order invariance), while the other axioms are very similar. The setup is almost identical as well. This time all individuals are placed into four, possibly overlapping, categories: X, $\neg X$, Y, $\neg Y$. Once again, dyads are chosen randomly without replacement until all $\frac{n(n-1)}{2}$ have been selected, and each dyad has to option to form a link. We are interested in measuring the proclivity of a member of X to a member of Y. In the discussion below, we will define $D_{i,j}$ to be the set of dyads, with some ordering, that can be constructed between member of groups i and j, so, for instance, $D_{X,Y}$ is the ordered set of dyads formed between members of X and Y. To simplify the notation in the following part, the following definition is useful.

Definition 4.4. *Let $R \subseteq D_{X,Y}$, and $S \subseteq D_{X,\neg Y}$ be the (ordered) subsets of the dyad sets that have links between them, then $D^*_{X,Y}(R,S)$ is a process associated with R and S, involving groups X and Y.*

The axioms are stated below to define a proclivity measure for X towards Y, g:

Two-Way Monotonicity. *More links to members of Y by X yields a higher value, and more links outside of Y leads to a lower value. Define $X, Y \subseteq N$. Let $A, B \subseteq D_{X,Y}$ and $A', B' \subseteq D_{X,\neg Y}$ be subsets of links between dyads. If $A \subseteq B$ and $B' \subsetneq A'$, or $A \subsetneq B$ and $B' \subseteq A'$, then $g(D^*(A, A')) < g(D^*_{X,Y}(B, B'))$.*

Group Information Only. *The index is only dependent upon group labels, so there is anonymity within groups. Let $R \subseteq D_{X,Y}$ and $S \subseteq D_{X,Y}$ be ordered subsets of the dyad sets. Furthermore, let $\Pi D_{X,Y}$ and $\Pi D_{X,\neg Y}$ be permutations of $D_{X,Y}$ and $D_{X,\neg Y}$, respectively. Then, $g(D^*_{\Pi X,\Pi Y}(R,S)) = g(D^*_{X,Y}(R,S))$.*

Order Invariance. *Dyads have same impact on the index, irrespective of when they are chosen. Formally, we say $g\left(\Pi D^*_{X,Y}\right) = g\left(D^*_{X,Y}\right)$ where $\Pi D^*_{X,Y}$ is a permutation of $D^*_{X,Y}$.*

Intuitively, the axioms are as they were in defining Z in the section above, but now individuals use group information to form links. The first condition just states that, all else being even, proclivity of X towards Y increases as more links are formed between members of X with members of Y and decreases, all else being even, when links are formed outside of Y. The second condition allows individuals to form links based upon group information only, and the order invariance is a condition carried over from before so that impact of forming a link on the proclivity measure is independent of when in the process it was formed. Taking these three axioms together, we can once again form our index of interest.

Theorem 4.5. *The proclivity index between X and Y, $Q_{X,Y}$, is the unique function, $g : P\left(D^*_{X,Y}\right) \to [\alpha, \beta]$ that satisfies two-way monotonicity, group information only, and order invariance with a range of $[-1, 1]$ with $g\left(D^*_{X,Y}(\emptyset,\emptyset)\right) = 0$ for a graph with no isolated vertices.*

Proof: Consider a dyad, k, involving a node from X, and some process $D^*_{X,Y}(R,S)$. If k is a dyad between members of X and Y and forms a link (when it didn't do so before), then the marginal contribution of the link formed by k is:

$$g\left(D^*_{X,Y}(R \cup \{k\}, S)\right) - g\left(D^*_{X,Y}(R,S)\right) = c_k^{X,Y}$$

Similarly, the marginal contribution of a node in X forming a link in $\neg Y$ is simply:

$$g\left(D^*_{X,Y}(R, S \cup \{k\})\right) - g\left(D^*_{X,Y}(R,S)\right) = c_k^{X,\neg Y}$$

As before, order invariance requires that $c_k^{X,Y}$ and $c_k^{X,\neg Y}$ be constants. Groups only requires anonymity within groups, so we can set $c_k^{X,Y} = \alpha_1$ and $c_k^{X,\neg Y} = \alpha_2$, where α_1 and α_2 are constants. It follows that g is "quasilinear " in the number of links between X and Y and between X and $\neg Y$. So:

$$g\left(D^*_{X,Y}\right) = \alpha_1 \times \sum_{k \in D_{X,Y}} d_k - \alpha_2 \times \sum_{j \in D_{X,\neg Y}} d_j + h\left(D^*_{\neg X, \neg X}\right)$$

where $d_k, d_j = 1$ if a link is formed in the dyad and h is some function of the links formed with both nodes in $\neg X$. Monotonicity implies $\alpha_1, \alpha_2 > 0$, and the condition $g\left(D^*_{X,Y}(\emptyset, \emptyset)\right) = 0$ implies $h \equiv 0$. Finally, a range of $[0,1]$ implies:

$$\alpha_1 = \frac{2}{n_Y(n_Y - 1)}; \quad \alpha_2 = \frac{2}{n_{\neg Y}(n_{\neg Y} - 1)}$$

Once again we use the fact that:

$$2 \times \sum_{k \in D_{I,J}} d_k = \sum_{i \in I} \delta^J$$

Substituting in directly gives us:

$$g = \frac{1}{n_Y(n_Y - 1)} \sum_{i \in X} \delta_i^Y - \frac{1}{n_{\neg Y}(n_{\neg Y} - 1)} \sum_{i \in X} \delta_i^{\neg Y} = Q_{X,Y} \quad \square$$

4.2 Stochastic Approaches to Proclivity

Once again, while the axiomatic approach gives a view of the proclivity measure as a deterministic quantity, the measure has a clear stochastic interpretation, the difference between the probability of a member of X forming a link with a member of Y and the probability of forming a link with $\neg Y$.

Since the choices of X and Y are up to the researcher, there is some concern about manipulation of the numbers. In particular, the estimates for proclivity might be unreliable in small groupings. As an extreme case, we can conceive of X and Y as two (distinct) nodes with a link between them. Do we really want to believe that the probability of forming a link between X and Y is 1? When thinking about the proclivity measure, the sample size concerns are more acute because we are subsetting the data. As before, I propose a Bayesian routine to deal with these concerns.

It is useful to understand proclivity in a regression context. In the frequentist context, we restrict our data to all of the dyads that have at least one member of X (i.e. we throw out the dyads in $D_{\neg X, \neg X}$). Let d_k be 1 if a link is formed in dyad k, and 0 otherwise. Furthermore, define $\omega_{X,Y}$ to be 1 if the dyad is in

$D_{X,Y}$ and 0 otherwise (i.e. the dyad is in $D_{X,\neg Y}$. We may quickly calculate $Q_{X,Y}$ by running the following regression:

$$logit(d_k) = \beta_0 + \beta_1 \times \omega_{X,Y}$$

where β_0 and β_1 are parameters estimated from the regression. We would then estimate:

$$Q_{X,Y} = logit^{-1}(\beta_0 + \beta_1) - logit^{-1}(\beta_0)$$

Notice that the first term on the right side is just the (frequentist) probability of forming a link in $D_{X,Y}$ and the second term is just the probability of forming a link in $D_{X,\neg Y}$. Thus, we can quickly calculate $Q_{X,Y}$ as a frequentist measure in a regression context.

We can also use a similar setup to derive a Bayesian estimate. The problem with the frequentist estimate is the researcher may choose small X and Y groupings, for which there is less confidence in the measure. In order to account for this situation, we run a random intercepts model, where the dyads are broken into three groups: $G = \{D_{X,Y}, D_{X,\neg Y}, D_{\neg X,\neg X}\}$.[8]. Formally, we run:

$$logit(d_k) = \alpha_i; \quad \alpha_i \sim N(\mu_\alpha, \sigma_\alpha)$$

where $i \in G$ and μ_α and σ_α are hyperparameters used to estimate the random effects model (i.e. the random coefficients are assumed to be normally distributed with some common mean and variance). In this case, the Bayesian estimate of proclivity of X towards Y is:

$$Q^B_{X,Y} = logit^{-1}\left(\alpha_{D_{X,Y}}\right) - logit^{-1}\left(\alpha_{D_{X,\neg Y}}\right)$$

It is advisable to run simulations to determine standard errors for each of these estimates. Since there is no "error term" in a binary logistic regression as there is in least squares, we need our calculate our standard errors off of the standard error estimates for each of the parameters in the regression models. Thus, in the frequentist case, we would use the standard errors on α and β to simulate the data-level standard error. In the Bayesian case, we would simulate off of the standard error estimates of each of the random intercepts. While these estimates do not have simple closed form values, modern computing technology makes this approach quick and straightforward.[9]

4.3 Defining Cleavage Structures

Naturally, we are interested in methods to select the groups; in other words, we want to endogenize the choice of X and Y in the proclivity measure. In order to endogenize the process, we introduce a mathematical notion of cleavage structure. When we think of a social/political cleavage, like an ethnic or linguistic cleavage, we tend to think of a set of "groups" that partition society[10] and have an aversion to each other (i.e. a strong proclivity towards their own group). Using this logic of a cleavage, and the proclivity measure discussed above, we construct endogenously defined groups. The following definition is useful for our purposes:

[8] We may, also, if we choose, further partition $D_{\neg X,\neg X}$ into dyads that have a link with Y and those that have a link with $\neg Y$, so that the inferences are more consistent with the latin square motivation.

[9] For a nice introduction to this approach, please see the discussion in Gelman and Hill (2006).

[10] There is some criticism as to whether group/cleavage structures should partition society. For instance, Airoldi, Blei, Fienberg and Xing (2008) argue that rather than viewing individuals as members of one group, they should be viewed as having a probability distribution over all groups in society. While appealing in many ways, this argument is more a criticism of the notion of cleavage structures rather than a definitional problem.

Definition 4.6. *A cleavage structure of order M is a partition of the non-empty graph with no isolated vertices, G, $\mathcal{P} = \{x_1, x_2, ..., x_J\}$ such that*

$$M = \sup\{K | Q_{x_i,x_i} > K \; \forall x_i, x_j \in \mathcal{P}, \; i \neq j\}$$

In other words, M, is the greatest lower bound of the proclivity measures for all of the groups in the partition. This definition of cleavage structures allows us to distinguish a cleavage-type structure from a process that might be better described as "ghettoization" rather than a true cleavage structure. In other words, the reciprocal aversion between groups in cleavages allows us to get high values of M. The following graphic shows the basic logic:

Figure 5: Comparison of Social Structures

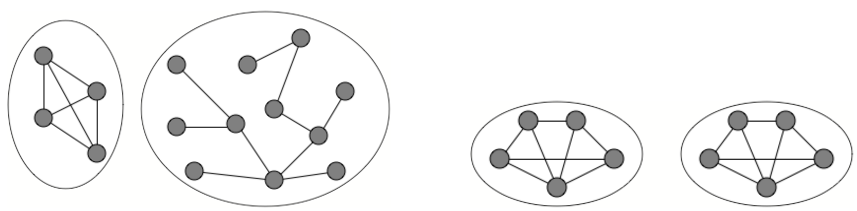

(a) Ghettoization: M=0.22 (b) Cleavage: M=1

The graphs above depict two groups under different social structures. Clearly, the structure in (b) admits a much higher value of M due to reciprocal aversion.

The value of M also rapidly decreases with the size of groups in the partition. The following graphic shows that it is very hard to have high values of M with partitions composed of small groups. Just by adding two links, we are able drop M by 19%. This also explains why this approach is not just a simple clustering algorithm. In some sense, the size of the groups is endogenously defined within this definition, and the cleavage structure biases towards larger groupings (which is consistent with our intuitions about social cleavages being composed of a few large groups).

This simple framework of cleavage structure, leads to an intuitive definition of the *most salient cleavage* in society:

Definition 4.7. *Let $\mathbf{P_G}$ be the collection of possible partitions of the graph G. The cleavage function, $\mathcal{C}_G : \mathbf{P_G} \rightarrow [-1, 1]$ takes each partition to the order (M) of the cleavage. The most salient cleavage is the partition:*

$$\mathcal{P}^* = \arg\max \; \mathcal{C}(\mathcal{P}), \; \mathcal{P} \in \mathbf{P_G}$$

This cleavage may not be unique. To see this, notice that any partition over a complete graph will yield $M = 0$ for all groups. Nonetheless, we can put a sharper lower bound on the cleavage order of the most salient cleavage:

Theorem 4.8. *Let $M^* = \mathcal{C}(\mathcal{P}^*)$; $M^* \in [0, 1]$ and $M^* = 0$ if and only if the graph, G is complete.*

Proof: M^* is bounded above by 1. If G is a complete graph, then clearly $M^* = 0$ since each partition of the graph yields $M = 0$. To prove the theorem, it suffices to show there exists a partition $\mathcal{P}'$ of the graph such that $\mathcal{C}_G(\mathcal{P}') > 0$ for all the G that are not complete.

Consider the coarsest partition of G such that each element of the partition is a complete subgraph. This is always feasible when the number vertices is even (a list of edges satisfies this condition). It is trivial to show that the proclivity measure is greater than 0 for each element of the partition. Now consider the scenario when the number of vertices is odd. Remove the lowest degree vertex (if the partition above is not feasible) and create a partition as above and then add the vertex to the largest component. Direct calculation shows that each proclivity measure is again positive. $\square$

Figure 6: Comparison of Small Group Structures

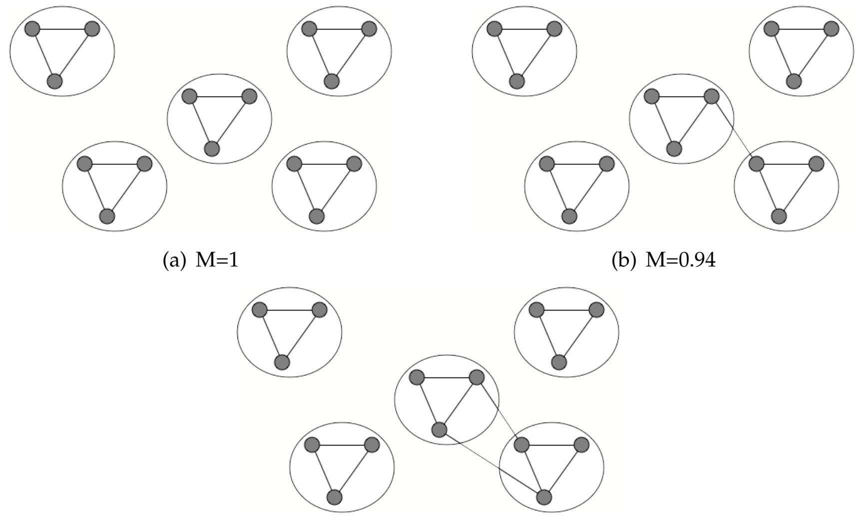

(a) M=1

(b) M=0.94

(c) M=0.81

The graphs above depict how M decreases rapidly as we add links. Between (a) and (c), M decreases 19% by adding 2 links. Also note that we can factor in the uncertainty due to the size of the groups by using Bayesian measures of proclivity, which would lower M even further for smaller groups.

5 Conclusion

This paper develops three measures for the study of social fragmentation using non-directed network data: fractionalization, proclivity, and most salient cleavage. Fractionalization measures the level of disconnectedness in the sample. Proclivity measures the preferences in forming a link with members of Y by members of X (as compared to those outside of Y). Finally, the most salient cleavage is a method to partition the network into meaningful clusters. In the paper, we begin by interpreting network data in terms of a theoretical framework where each individual has the opportunity to interact with every other individual. Using this approach, we develop measures of fractionalization and proclivity measures from both an axiomatic and statistical perspective.

The data demands for the type of analysis described in the paper are quite high. At the very least, it represents an ideal scenario for estimation of social fragmentation. The measures can be adapted to more easily obtainable data, like that of directed networks which can often be obtained by surveys. Furthermore, it may be possible to obtain non-directed network data which conform to the assumptions of the paper in an experimental setting or some micro-level structure (e.g. a school or an apartment building). This paper represents a first step in developing a framework where social fragmentation can be estimated more accurately from network-type data. Hopefully, further research will help determine large-scale estimation techniques that can approximate the structure in this paper.

References

Airoldi, E.M., D.M. Blei, S.E. Fienberg and E.P. Xing. 2008. "Mixed membership stochastic models." *Journal of Machine Learning Research* 9(1981-2014).

Atlas Narodov Mira. 1964. Moscow: Miklukho-Maklai Ethnological Institute at the Department of Geodesy and Cartography of the State Geological Committee of the Soviet Union.

Esteban, Joan-Maria and Debraj Ray. 1994. "On the Measurement of Polarization." *Econometrica* 62(4):819–851.

Fearon, James. 2003. "Ethnic Structure and Cultural Diversity by Country." *Journal of Economic Growth* 8(June):195–222.

Gelman, Andrew and Jennifer Hill. 2006. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press.

Gelman, Andrew, John B. Carlin, Hal S. Stern and Donald B. Rubin. 2003. *Bayesian Data Analysis*. Chapman and Hall/CRC.

Goffman, Erving. 1983. "The Interaction Order: American Sociological Association, 1982 Presidential Address." *American Sociological Review* 48(1):1–17.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78(6):1360–1380.

Jackson, Matthew O. 2008. *Social and Economic Networks*. Princeton University Press.

Krackhardt, David. 1994. Graph theoretical dimensions of informal organizations. In *Computational Organizational Theory*, ed. Kathleen Carley and Michael Prietula. Lawrence Erlbaum Associates, Inc.

Laitin, David and Daniel N. Posner. 2001. "The Implications of Constructivism for Constructing Ethnic Fractionalization Indices." *APSA-CP* 12:13–17.

Linz, Juan and Alfred Stepan. 1996. *Problems of Democratic Transition and Consolidation*. Johns Hopkins University Press.

Posner, Daniel N. 2004. "Measuring Ethnic Fractionalization in Africa." *American Journal of Political Science* 48:849–863.

Sen, Amartya. 2006. "What Clash of Civilizations?" Slate (March 26, 2006).
URL: *http://www.slate.com/id/2138731/*

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