

Simulating and Analyzing Dark Networks: Modeling and Measuring Using Network Tools

David C. (Chris) Arney, Jocelyn R. Bell, Kathryn A. Coronges,
& Greg Merkl

The debate over the relevance of network science in effectively analyzing dark networks continues for two reasons: the method's historical record is short in duration and the results are mixed (Alberts & Hayes, 2005; Xu & Chen, 2008; Jones, 2012). The few trusted data sets and formal assessments that do exist remain contentious. Clandestine organizations deliberately guard their structure and process so that attempts to objectively categorize their members and their relationships invite criticism and scrutiny. Roberts and Everton (2011) conducted a deep analysis of Jemaah Islamiyah in an attempt to provide an accepted case study of the formation of a terrorist organization. Similarly, Krebs (2002a, 2002b) categorized al-Qaida following the September 11 attacks. Both works add meaningful data to the historical record, but neither paper comes without controversy as to the accuracy and completeness of its analysis.

The debate pitches the proponents of dark network analysis against skeptics who doubt the utility of simplifying the complexity of real-world clandestine relationships into formal representations of nodes, edges, processes, and attributes of a network graph. Several case studies give merit to both sides of the debate (Arreguin-Toft, 2001; Philby, 2013). On the micro scale of dark network analysis, centrality metrics can offer meaningful insights into the structure of subgroups and the power players within the organization (Borgatti, 2006; Borgatti, Carley, & Krackhardt, 2006). However, on the macro scale, the changing nature of information available to analysts can cause confirmation biases. That is, traditional full-graph measures can be biased both by data availability and the analyst's erroneous belief about what data is missing. Therefore, depending on data reliability, traditional centrality measures of the known elements

The views expressed herein are those of the authors and do not purport to represent the official policy or position of the United States Military Academy, the Department of the Army, the Department of Defense, or the U.S. government.

of a dark network may provide little help to understand the real unknown dark network.

The dramatic dynamics associated with some dark networks present an additional complication. Elements and agents can coalesce and metastasize to produce unpredictable events, appear in unanticipated places, and create numerous paradigm shifts producing significant second- and third-order effects. These factors make dark network analysis a challenging venture that is much like Silver's (2012) description of finding the signal through the noise. In this case, the signal consists of the real elements of the dark network, while the noise is all the secrecy, complexity, and dynamics that create a hidden shield around the dark organization.

Even though these challenges are significant, the reward of trustworthy data collection and sound network analysis can be substantial. There are many military threats that involve dark networks, and by some measures, all enemy organizational networks are dark (Galula, 2006; Gartenstein-Ross, 2011). A military force never wants its opponent to know about its organization or operations (Flynn, Pottinger, & Batchelor, 2010). The following military operational threats often involve network-based components or utilize network models within the intelligence-gathering operations:

- *Irregular warfare* – Understand the employment of unconventional, asymmetric methods to include terrorism, insurgency, and guerrilla warfare.
- *Regular warfare* – Intelligence determination of enemy order of battle and critical enemy units such as special operations cells that work with the local populace.
- *Disruptive operations* – Understand the network-enhanced technologies that reduce the U.S. advantage in cyber operational domains such as identifying the cyber cell that is conducting offensive cyber operations.

Recent counterinsurgency operations have definitely featured the valuable role of network analysis in many levels of intelligence gathering (Bolz, Dudonis, & Schulz, 2002; Sageman, 2004). As described in the U.S. Army's field manual on counterinsurgency operations, the military is using social network analysis in its intelligence strategy:

“Social network analysis (SNA) is a tool for understanding the organizational dynamics of an insurgency and how best to attack or exploit it. It allows analysts to identify and portray the details of a network structure. It shows how an insurgency's networked organization behaves and how that connectivity affects its behavior. SNA allows analysts to assess the network's design, how its members may or may not act autonomously, where the leadership resides or how it is distributed among members, and

how hierarchical dynamics may mix or not mix with network dynamics.” (FM 3–24, p. B-10)

We need more realistic and accurate assumptions to inform models and analyses of dark networks. A viable set of assumptions could be:

- Actors in a dark network sometimes act in irrational or chaotic ways.
- Data and information for the nodes, links, and their attributes may be inaccurate, unknown, or missing.
- Times-series situation is dated and could be inaccurate.
- Structure is variable and often evolves in unpredictable ways.
- Processes and procedures for dark networks are not consistent.
- Layers and dimensions of connections are complex (as shown in Figure 8.3).

Despite these forceful assumptions, nonreductive network modeling is still the most powerful method we have for dark networks. As this volume indicates, tools and methods are being developed to meet these kinds of challenges in dark network modeling, analysis, and synthesis (Brandes et al., 2013). The techniques we show in this chapter contribute to advancing these developments.

Our framework for dark networks attempts to address the fundamental questions under study: *How do dark networks operate? How are they structured? Where are they vulnerable?* Our analytic approach embraces the complexity of systems, reveals and synthesizes their complex structures and processes, and provides usable metrics and models.

From a military perspective, the objective we assume is to effectively attack the network. Specifically, we seek to identify the most important and/or weakest agents, name the most powerful or vulnerable subgroups, and build a viable target and attack methodology. The following two examples exhibit the methodology and utility of these new tools.

I. Bell’s Subgroup Technique: Identifying Hidden Targets

In certain situations, a dark network may be embedded in a light network. In other scenarios, there may be certain individuals within a dark network who cannot be targeted for removal because of diplomatic concerns or other political consequences. In these cases, we must decide which of the targetable individuals should be removed to cause the most disruption. To do this, we need to know which individuals have the most influence over the targetable portion of the network.

Existing approaches to analyze this network would be to either: 1) consider measures on the targeted-only network, or 2) consider measures

on the combined network. Because some network connections are either ignored (as in case 1) or treated as the same type of connections (as in case 2), neither of these alternatives is accurate or meaningful in this situation. Completely ignoring the legitimate connections in a dark network may be ill advised, but these connections are different in purpose than the illegal connections within the targeted network. Thus, traditional centrality measures fall short of answering important questions about network structure. The subgroup measures of Bell (2014) take both micro- and macro-level settings into account by allowing for the division of the network into local (targeted) and global (untargeted) influence. Bell's subgroup technique generates centrality rankings that differ substantially from the traditional approaches.

Bell (2014) presented a new framework where nodal centralities are calculated accounting for both local structure (within the subgroup) and global measures (the entire network outside the subgroup). Normalized measures take the size of the subgroup under investigation into account. A subgroup measure according to subgroup S is defined as:

$$a \in V \text{ (the set of vertices),}$$

$$C_S = \sum_{x \in S} f(a, x)$$

The sum is restricted to only those nodes in S , and f is the value of the relationship of its two nodal elements (a and x) to the centrality measure. We say a subset S of the set of vertices V is a subgroup of the network; note this is not the subgraph induced by S as it contains no edge information. There are two special cases when a happens to be an element of S . The subgroup measure of a according to S is a local measure, which measures how central node a is inside the subgroup. The subgroup measure of a according to S^c (everything except the nodes in S) is a global measure, which measures the centrality of a with respect to nodes outside the subgroup. According to these definitions, the sum of a node's local measure and its global measure equals the value of the original measure.

We demonstrate by means of an example how Bell's subgroup technique is more effective in identifying and targeting key targetable nodes embedded within larger non-covert networks. To demonstrate the performance of the subgroup-based centrality measures, we will use the classic dolphin relationship network presented by Lusseau and others (Lusseau et al., 2003). For the purpose of this demonstration, we will consider the targetable population of thirty-three individuals, an untargetable population of twenty-five, and four individuals of unknown status. Figure 8.1 depicts the entire network, whereas Figure 8.2 shows the network of targetable individuals and ignores connections through the legal network.

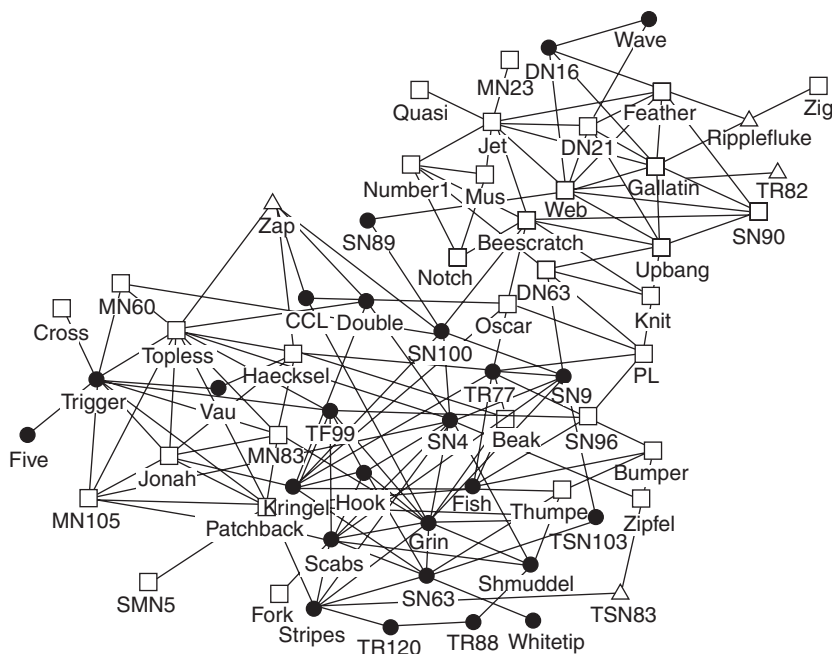


Figure 8.1. Entwined network of targeted (squares), untargeted (disks), and undetermined (triangles) populations.

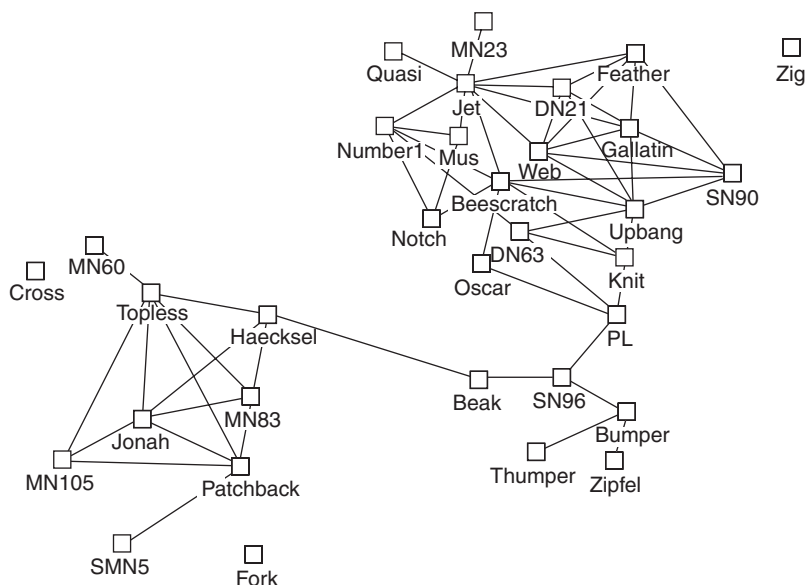


Figure 8.2. Targeted network only.

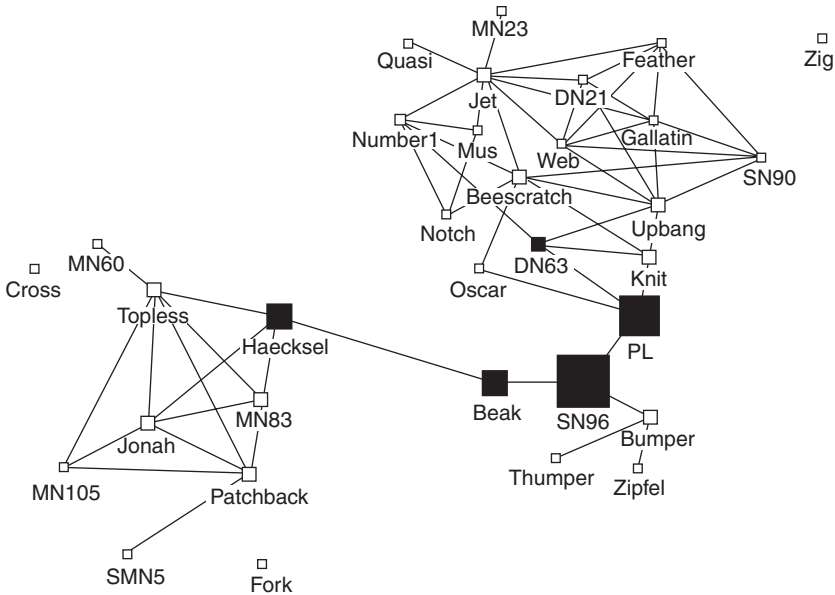


Figure 8.3. Agents sized by traditional betweenness.

We compare the traditional approach (calculating measures on the targetable network only) versus Bell's subgroup-based measures using betweenness, a measure of centrality useful for determining key nodes in network information flow (Freeman, 1977). In Figure 8.3, betweenness has been calculated in the usual sense on the targetable network; agents are sized according to their betweenness, and the five most central individuals are represented by solid squares. Although SN96 and "Beak" are among the most central nodes, their removal will not prevent the resulting network fragments from communicating, as messages may still travel through the untargetable population. Figure 8.4, which depicts Bell's local betweenness measure but otherwise follows the same visualization conventions as the previous graphic, reveals a vastly different power ranking, suggesting that influencing Beescratch rather than SN96 will have the most impact on disrupting communication among the targetable population.

In other cases, we may wish to disrupt or influence the untargetable population, but are restricted to removing, disrupting, or influencing targetable agents only. Bell's global version of these subgroup measures accomplishes this by quantifying the influence a particular node has outside their subgroup. Figures 8.5 and 8.6, which again size and color nodes based on centrality, compare Bell's local and global closeness measures, indicating key positions in either the targetable population (Figure 8.5) or the untargetable population (Figure 8.6). To damage the untargetable

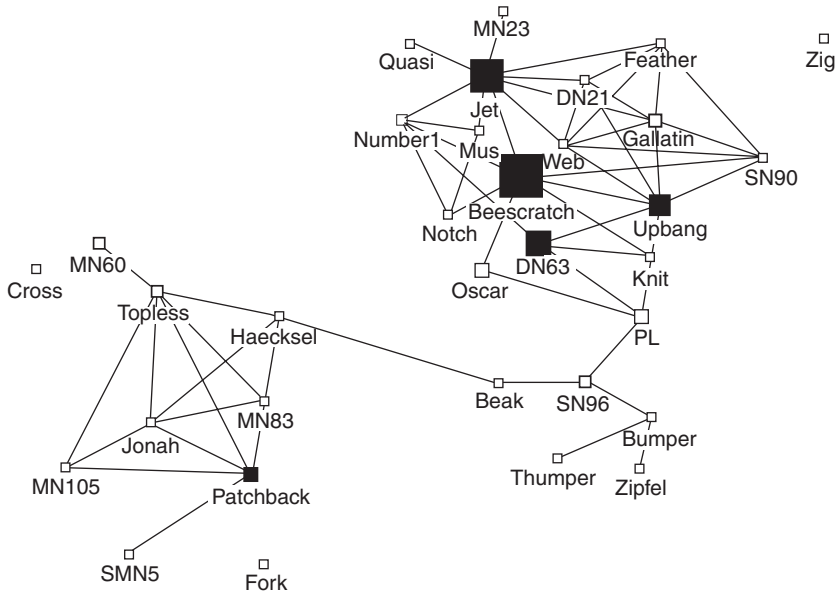


Figure 8.4. Agents sized by Bell's subgroup betweenness.

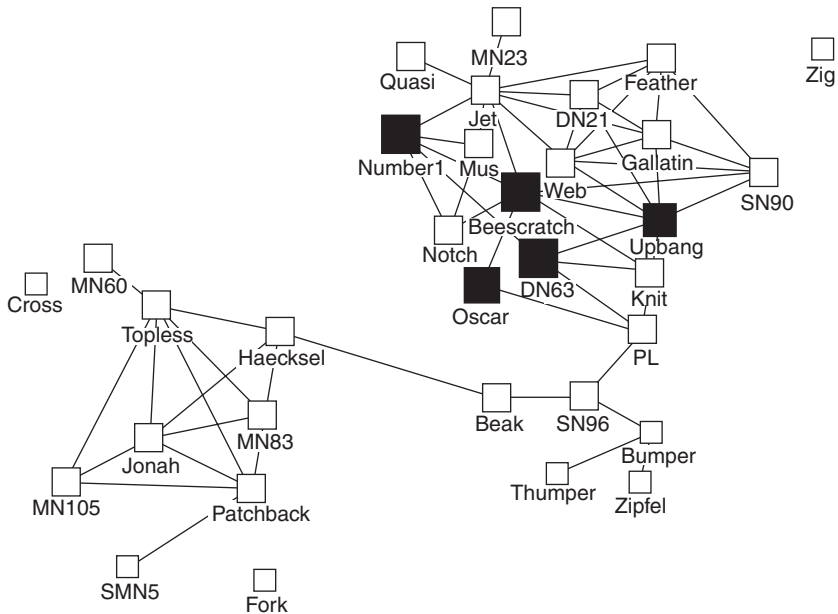


Figure 8.5. Nodes sized according to Bell's local closeness.

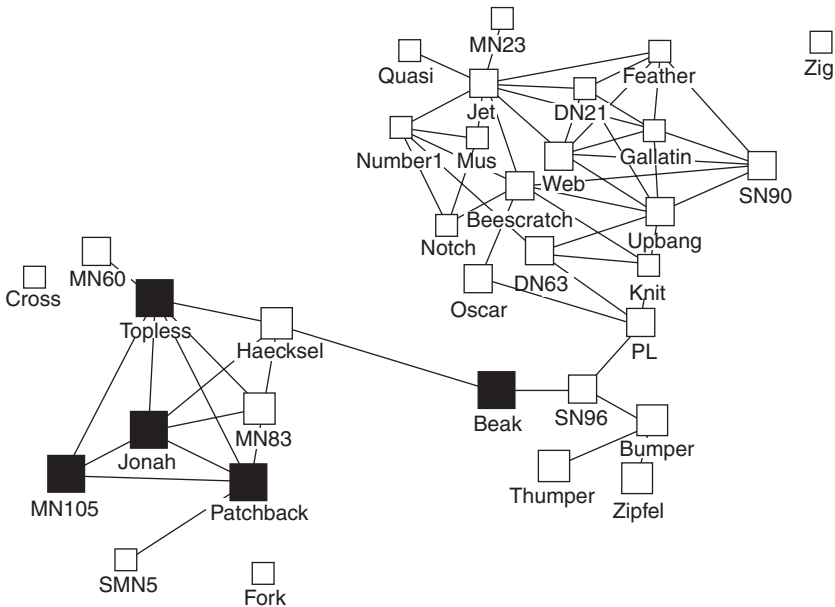


Figure 8.6. Nodes sized according to Bell's global closeness.

population, Bell's method suggests we remove Topless, while targeting Beescratch damages the targetable population.

A third use for Bell's approach is to help determine which classification the undetermined individuals should belong to. If an undetermined node has more local (targetable) influence than global (untargetable), it suggests that this individual should be classified as targetable. Unknown individuals are linked with the targetable population in Figure 8.7 and with the untargetable population in Figure 8.8. Ripplefluke and TR82 probably do not belong in the untargeted group, as their only connections to the network are to targetable members.

Other individuals are not as clear cut. Consider TSN83. His global closeness rank (18th) is significantly higher than his local closeness rank (34th). We might conclude that TSN83 has more influence over the untargetable population and so place him in that category. Zap is an even tougher case; he is well connected to both groups. His power rankings in both local and global closeness are similar (10th and 9th, respectively). However, there is a significant rank difference for Zap in local betweenness (21st) versus global betweenness (11th). This may persuade us to classify Zap as targetable.

Further discovery can be made from comparing the raw (un-normalized) global centrality score to the raw internal score. While this does not enable ranking individuals, it does indicate to which subgroup the individual is more strongly attached. E-I (External-Internal) index measure

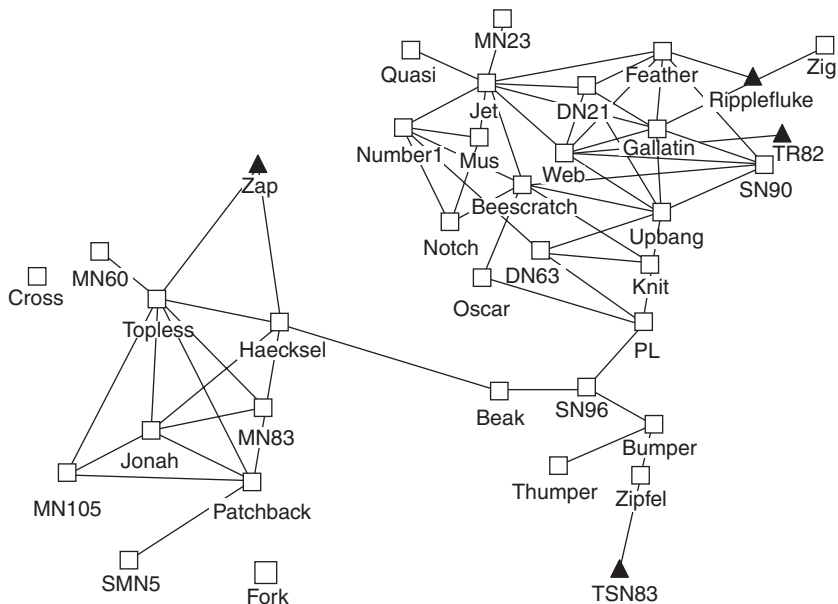


Figure 8.7. Unknown individuals (triangles) and targetable population (squares).




Figure 8.8. Unknown individuals (triangles) and untarable population (circles).

Table 8.1. *Global-local (E-I) indexes using Bell's subgroup-based centrality measures*

	Degree E-I Index	Closeness E-I Index	Eigenvector E-I Index	Betweenness E-I index
Ripplefluke	-1	-0.05674	-1	-0.16667
TR82	-1	-0.09924	-1	0
TSN83	0	-0.32766	0.570905	0.633588
Zap	0.2	-0.24719	-0.11684	0.434212

(Everett & Borgatti, 2012) can be calculated where E = un-normalized global score and I = un-normalized local score.

$$E - I \text{ Index} = (E - I) / (E + I)$$

A positive E-I index indicates a stronger tie to the untargetable portion of the network than the targetable. The E-I indexes for the unknown individuals appear in Table 8.1.

The E-I indices largely support Bell's subgroup analysis approach; Ripplefluke and TR82 are more likely to belong to the targetable group, while TSN83 appears to belong to the untargetable group. However, there is an even split for Zap, so here the E-I index combined with Bell's method may be the best approach.

While Bell's approach offers means to infer group membership, this method cannot account for network errors stemming from one of the most challenging aspects of data collection in dark networks: link discovery. Consequently, the next section turns its attention toward Merkl's sampling technique, which realistically simulates the data discovery process undertaken by human analysts. Given that clandestine organizations attempt to obscure their membership, making it impossible to know the true size of a dark network, his approach offers one means to determine the fidelity of analytic conclusions when analysts lack basic demographic information about the networks under study.

II. Merkl's Dark Sampling Technique: A Bayesian Methodology to Simulate Network Evolution

In this example, we synthesize a dark network based on probabilistic elements within the uncovered network. The underlying assumptions are that clandestine organizations tend to have similar structure, and that analysts tend to discover elements of dark networks via well-defined, but

inherently probabilistic processes (Flynn et al., 2010; Hung, Kolitz, & Ozdaglen, 2011). Our targeting algorithm simulates the workflow analysts typically undertake as they probe the boundaries and membership of an illicit organization; the simulation uncovers clandestine networks using a Bayesian approach and changes strategy as more information is acquired. The combination of these two conditions yields a dark network discovery method to uncover the edges of a dark network that, like Bell's work, affirms the validity of traditional centrality measures on the micro (subgroup) level while simultaneously questioning their relevance at the macro level.

The simulation assumes that analysts start with an equal probability of discovering all edges. Then, after receiving information about the connected vertices and the changing security environment, analysts have a greater propensity to discover vertices connected to previously discovered vertices. Thus, the simulation iteratively updates nodal and link assumptions based on previously known information.

Analysts seeking to discover the actors within a dark network must make economic decisions about how to dedicate limited intelligence resources. As a general rule, once analysts have uncovered an actor within a network, they have a greater propensity to discover the actor's close associates than distant members of the network. This pattern results in a loose version of snowball sampling (Goodman, 1961). However, we utilize this snowball sampling procedure as it pertains specifically to dark networks by formalizing the probabilistic method of discovery. We describe the three steps and corresponding assumptions in what we call *dark sampling*.

- *Uncover the edges* – Analysts seek to uncover the edges (links) between vertices (actors) within a dark network. Often, the actors lead double lives that are transparent in one setting, but that disappear into the shadows of secrecy in another. Consider the September 11 bombers; they had legitimate paperwork, conducted economic activities under their true identities, and lived among the community. They were clearly visible to the U.S. government. Their connections to al-Qaida, however, remained carefully guarded (Kean & Hamilton, 2004). Consequently, our simulation focuses on uncovering the edges of the network in order to reveal the nodes within the organization.
- *Take advantage of non-secret edges and vertices* – Analysts struggle to discover actors within the network because they maintain some level of secrecy. However, a lack of operational security among any one member of the network will unilaterally degrade the security of his neighbors. This condition frequently occurs in real-world situations, and for this reason some of America's

greatest missions against al-Qaida have started by tracking a courier to the location of a commander (Bowden, 2012). To represent this mathematically, we assume that each node in the network has a level of secrecy that ranges between 0 and 1. For instance, assume that Alice and Bob have joined al-Qaida and Alice maintains fierce vigilance with a high secrecy level of 0.99. Bob, however, tends to suffer from loose lips as he talks to others on his way home from the office. His secrecy level is 0.25. Assuming independence of the probabilities, the chance that analysts will fail to discover the link between Alice and Bob is $(0.99)(0.25) = 0.2475$. To account for the cyclic nature of targeting, we assume that these probabilities are based on some time unit. For instance, for a given week, our simulation has a 75.25 percent chance of discovering the connection between Alice and Bob.

- *Illuminate hidden regions of the network* – Once analysts discover a node, its secrecy suffers. Real-world targeting strategies tend to use these kinds of Bayesian methodologies. Once analysts locate Alice and Bob as members of the dark network, their friends and neighbors have a higher propensity to share affiliations with the same network. Our simulation consequently diverts resources toward Alice and Bob, thereby reducing their ability to maintain hidden relationships and lowering their security level in our algorithm.

Figure 8.9 depicts a simulation of the dark network discovery process described earlier. The network under study is sparse and wide in diameter, highlighting the potential for the simulation to discover multiple components of a connected graph. The figure's left-hand panel shows the actual complete network – perfect information that analysts are unlikely to ever entirely discover. The figure's right-hand panel depicts a partial sampling of the edges after discovering approximately 25 percent of the edges using our methodology. In both panels, the most central actors are represented by darkly shaded nodes; thus, readers can intuitively see differences between the actual network and the “Targeter View” uncovered by analysts by noting that the dark clusters are located in different sections of the respective graphs.

More scientific performance estimates are available from the graph across the bottom of the figure, which depicts eigenvector centrality error as a function of the simulated analysts' knowledge of the network. The thick, volatile line depicts Spearman's Rho rank-order error (Spearman, 1904), while the thinner line hugging the origin of the graph shows the relative change in the centrality measure as the simulation obtains each piece of information about the network. The straight diagonal line shows a linear convergence rate between information and error; analysts seek




Figure 8.9. Simulating the discovery of 25 percent of links in a dark network.

to stay below the line, where the error rate is lower than the available information.

For this simulation, the hypothetical analysts made several radical changes in their centrality calculations while discovering the first 5 percent of the network. Initially, the new information aided analysts in correctly rank ordering the importance of the actors according to eigenvector centrality. However, shortly after obtaining 5 percent of the possible information, the analysts' new information led them to draw increasingly bad conclusions regarding the composition of the network's core. As the amount of discovered information increases from 5 percent to 25 percent, the simulated analysts' error for rank ordering the central actors remains about 75 percent. This example demonstrates that simply discovering more information about a dark network does not always improve analytic conclusions; when a large percentage of the information about a dark network remains unknown, adding new data can degrade analytic accuracy.

Figure 8.9 shows another important aspect of dark sampling. While the information tends to spiral out from known actors, the probability of discovering actors not connected to the known network is nonzero. Therefore, analysts may initially discover the network as separate components. In Figure 8.9's right-hand panel, the simulated analysts see two separated components of the network. Information will tend to snowball from these two components, but as greater intelligence about the network becomes available, it could develop even more separated components. The potential for the discovery of multiple components constitutes a primary distinction between dark sampling and snowball sampling.




Figure 8.10. Simulating the discovery of 50 percent of links in a dark network.

Figure 8.10 revisits the same network, albeit after discovering a total of 50 percent of the links in the network. Because this graphic follows the same interpretive and visual conventions as Figure 8.9, it demonstrates the effects of confirmation bias. The simulated analysts have uncovered enough information to join the two components they knew at the earlier discovery point, and are now more certain that the darkly colored nodes in the lower-right quadrant of the right-hand panel represent central actors in the dark network. However, comparison with the true network depicted in the left-hand panel shows that the analysts are mistaken. The Spearman's Rho information in the graph at the bottom of the figure confirms the analysts' error; values have strayed above the convergence line, indicating that the amount of error in the network is greater than the available information. With 50 percent of the total information, analysts have nearly 70 percent error in rank ordering the actors according to eigenvector centrality. Increased intelligence has ironically led to a corresponding increase in error.

Fortunately, this pattern is not durable and reverses with the addition of another 20 percent of information. Figure 8.11, which depicts the discovery of 70 percent of all ties, shows an increasingly correct intelligence estimate of the network based on eigenvector centrality. The error rate for the "Targeter View" is now just 30 percent, indicating that the simulated analysts have achieved linear convergence. Their conclusions are as accurate as we would expect given the amount of information at their disposal, and a quick visual assessment of the darkly colored core nodes suggests that analysts are beginning to correctly identify the network's core.




Figure 8.11. Simulating the discovery of 70 percent of links in a dark network.




Figure 8.12. Simulating the discovery of 80 percent of links in a dark network.

However, analysts' success proves ephemeral. Figure 8.12, which shows the test network after discovering an additional 10 percent of information, depicts substantial analytic error. With a total of 80 percent of the links discovered, analysts have incorrectly identified the cluster at the lower left of the "Targeter View" as the central core of the network. The error graph at the bottom of the figure accordingly shows a worrying spike above the convergence line; error rates have returned to approximately 75 percent. The simulated analysts are again performing worse with more information.




Figure 8.13. Simulating the discovery of 100 percent of links in a dark network.

Figure 8.13 depicts a final end-state, in which analysts have discovered all possible links in the dark network. The graphic confirms that the volatility in accuracy seen at lower levels of link discovery persist until a surprisingly high percentage of the network is known. Although error rates plummeted as the rate of data discovery entered the mid-eighties, error rates spiked twice in the upper eighties, before finally reaching a steady convergence with the rate of available information once approximately 90 percent of ties were discovered. Thus, these results suggest that analytic conclusions derived from network science remain exceptionally sensitive to relatively small amounts of missing information. Considering that analysts are unlikely to ever discover all ties within a dark network, these simulations suggest that the potential for confirmation bias with the emergence of new information remains high.

These conclusions differ from the findings of studies that utilized other sampling methods. For example, Costenbade and Valente (2003) found that random sampling techniques, such as bootstrapping, produced centrality measurements that tended to converge quickly. The next section, therefore, compares our snowball-like approach to data discovery to a random approach.

III. Dark Sampling Implications for Subgroup Centrality

As a test case, we examined the same dolphin network used to illustrate Bell's approach to centrality. Because we seek to characterize the typical




Figure 8.14. Eigenvector centrality convergence for dolphin network.

performance of eigenvector centrality (the most stable of the canonical measurements of node-level influence and importance), we assessed aggregate performance of 1,000 discovery permutations. Figure 8.14 displays the results.

These results suggest that, on average, eigenvector centrality converges at a better-than-linear rate. However, note that the deviation in the error measure actually increases as the amount of available information increases from 10 percent to about 45 percent. This trend illustrates the confirmation bias. If analysts focus their attention on the core of the network early in the discovery process, eigenvector centrality measures will correctly rank order the actors. Conversely, if analysts wait until information on the periphery of the network is uncovered, the error in rank ordering will tend to increase with time, as the confirmation bias leads them away from the truth – until the real core is finally discovered.

The question remains: How does this pattern compare to random sampling approaches? We investigated this question by comparing permutations of our approach to data discovery with an equal number of random samples. Figure 8.15 displays the results and shows that the two approaches perform similarly, in terms of both convergence and variance. However, neither our approach nor random sampling produce particularly satisfying results; analysts would prefer to see decreasing variance with increasing information, a pattern that would indicate that network-based conclusions became more certain with the addition of new relational data.

Given the similar performance of these two approaches to sampling, it is not immediately evident that our approach offers any additional insight into “good” metrics given a partial information set. However, when we




Figure 8.15. Convergence and deviation of eigenvector centrality using two sampling approaches.

look at the change in measurement error as the targeter develops information, we see that dark sampling reduces error by an order of magnitude compared to random sampling. Figure 8.16 displays results for the same tests described in the preceding graphic and indicates that our snowball-like approach performs better on average. We believe that this advantage stems from the potential to discover the core early in the data-sampling process. When our approach finds a member of the core early in the data-discovery process, other members of the core will also be uncovered rapidly, and analytic conclusions about the most central actors in the network will remain relatively stable and largely accurate. However, when a random approach finds a member of the core early in the data discovery process, other members of the core are no more likely than peripheral nodes to be discovered next. Nothing ensures that the sample will latch onto the core and stay there, causing analytic accuracy to suffer accordingly.

IV. Conclusions

Our research suggests that while snowball-like approaches to data discovery are more advantageous to analysts than random approaches to data




Figure 8.16. Changes in eigenvector centrality using two sampling approaches.

discovery, caution is always merited when assessing incomplete information. Until analysts have discovered upward of 90 percent of the available links in a network, centrality rankings can fluctuate wildly, even for the most stable measure of node-level importance and influence. Analysts should expect that local centrality measures perform more accurately than global measures that seek to quantify the relative importance of every actor in the network. Consequently, analysts should consider using alternate measurements of centrality. Because Bell's approach to centrality can serve to contrast local importance with global importance, this method may provide one means for analysts to determine if the apparent centrality of an individual in the discovered network stems from his or her location in a small clique-like subgroup, or from true global importance to the network.

More broadly, this work indicates that complexity in its many forms, including nonlinearity, chaos, and randomness, permeates dark networks, rendering both traditional statistical regression and standard network metrics of limited utility (West & Grigolini, 2011). More powerful and sophisticated tools are necessary to understand the complexities of dark networks, and future research should investigate questions such as: How do new subgroup-based network measures change as legitimate nodes become part of the dark network and vice versa? How stable are

these measures as new members and connections are discovered? Do technology-based networks, such as those formed and mobilized through social media, operate and reveal themselves in similar fashion as face-to-face networks? How can complex multilayered network models further help targeting and intelligence processing in military and criminal operations (Thomas, Kiser, & Casebeer, 2005; Hampson, 2012; Singer, 2012; Kivela et al., 2014)? In the end, our research raises as many questions about dark network analysis as it answers.