The Manipulability of Centrality Measures —An Axiomatic Approach

Tomasz Wąs
University of Warsaw
Warsaw, Poland
t.was@mimuw.edu.pl

Marcin Waniek
New York University Abu Dhabi
Abu Dhabi, UAE
University of Warsaw
Warsaw, Poland
mjwaniek@gmail.com

Talal Rahwan
New York University Abu Dhabi
Abu Dhabi, UAE
trahwan@gmail.com

Tomasz Michalak
University of Warsaw
Warsaw, Poland
tpm@mimuw.edu.pl

ABSTRACT

Centrality measures are among the most fundamental tools for social network analysis. Since network data is often incomplete, erroneous, or otherwise manipulated, increasing attention has recently been paid to studying the sensitivity of centrality measures to such distortions. However, thus far no universal method of quantifying the manipulability of centrality measures has been proposed. To bridge this gap in the literature, we take an axiomatic approach. In particular, we introduce a set of intuitive axioms that characterize such a measure, and prove that there exists only one solution that satisfies them. Next, building upon this measure, we quantify the manipulability of the most fundamental centrality measures.

KEYWORDS

Centrality Measures; Networks; Manipulability

ACM Reference Format:
Tomasz Wąs, Marcin Waniek, Talal Rahwan, and Tomasz Michalak. 2020. The Manipulability of Centrality Measures —An Axiomatic Approach. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 9 pages.

1 INTRODUCTION

Centrality measures—methods for evaluating the nodes' relative positions and roles in the network—are among the most fundamental tools in social network analysis [25]. One of the issues that attracted attention in the literature is the *sensitivity* of centrality measures [8, 13]. This interest is driven by the fact that real-life data about the links in a network are often incomplete, erroneous, or otherwise distorted [17]. There are various reasons behind this, many of which are unintentional, such as the under-reporting of network relationships [28] (e.g., there are many real-life relationships that are not declared on Facebook) or the errors made by informants while asked about their ties [15]. The studies that evaluate the effects of such random distortions typically assume that only a certain percentage of the (randomly selected) links are known [6, 16, 19, 21] or there is a noise affecting the weights of the edges [26], and the analysis focuses on how the centrality-based ranking in such an incomplete or noisy network differs from the true one.

However, the sensitivity analysis based on random distortions is inadequate in situations where changes to the network do not occur by chance but rather as a a result of informed, rational decisions, i.e., due to *manipulation.* Since a straightforward *modus operandi* is to create fake accounts and/or add fake connections to boost the importance of certain network members and/or diminish others, various forms and magnitudes of manipulation are common in social networks [7, 12]. As a result, the interest in understanding how centrality measures can be manipulated has been recently growing in the literature. In particular, Crescenzi *et al.* [9] studied the problem of maximising Closeness centrality of a node by creating a limited amount of new edges incident to it. Analogous problems were also considered for Betweenness centrality [4], eccentricity [10, 23], and page-rank centralities [1, 22]. Also, Waniek *et al.* [30, 31] studied how an “evader” node could rewire a given number of edges in order to decrease her centrality.

Since in most cases considered in the above literature obtaining an optimal solution turned out to be intractable, a typical approach was to develop a heuristic as opposed to an exact algorithm. The manipulability of a given centrality was then studied by comparing the ranking of nodes before and after applying the heuristic. Hence, manipulability was assessed in the context of a particular heuristic and a particular centrality measure, which often precluded comparison of the manipulability of different measures.

In this paper, we take a more general approach, where we pose the question about *the theoretical underpinnings behind quantifying the manipulability of centrality measures.* To answer this question, we take an *axiomatic approach* and formulate the problem characterized by a network, an evader node, a centrality measure and a set of allowed actions. We then introduce seven axioms that we believe are reasonable requirements for a measure of manipulability. We then prove that there exists only a single measure that satisfies all of them. We call it the Average Minimal Actions Required (AMAR) measure as it is equal to the inverse of the minimum number of actions that must be taken to manipulate the position of the

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

evader node in the ranking averaged over all networks. We then use AMAR to experimentally quantify the manipulability of the four most popular centrality measures—Degree, Closeness, Betweenness and Eigenvector.

2 PRELIMINARIES

Graphs: We consider simple, undirected graphs. Every such a graph is a pair, $G = (V, E)$, where V is the set of nodes and E (or $E[G]$) is the set of edges, i.e.: $E \subseteq \{e \subseteq V : |e| = 2\}$. By $\mathbb{G}^V$ we denote the set of all possible graphs with nodes V. The set of *neighbours* of a node, $v \in V$, consists of all nodes connected to v via an edge, i.e., $\mathcal{N}_G(v) = \{u : \{u, v\} \in E[G]\}$.

A *path* is defined as a sequence of pairwise distinct nodes, $p = (v_1, \dots, v_k)$, such that each pair of consecutive nodes is connected by an edge, i.e., $\{v_i, v_{i+1}\} \in E[G]$ for every $i \in \{1, \dots, k-1\}$. The *length* of a path is the number of nodes that form it minus 1. For two nodes, $u, v \in V$, the *distance* between them, $dist_G(u, v)$ is defined as the minimal length of a path that starts in one of them and end in the other, or infinity if there is no such path. By $\Pi_G(u, v)$ we denote the set of all shortest paths between nodes u and v.

Centrality measures: *Centrality measure*, $F : \mathbb{G}^V \to \mathbb{R}^V_{\geq 0}$, is a function that for a given graph returns real nonnegative values reflecting the importance of the graph's nodes. In this paper, we consider four fundamental centrality measures:

- *Degree centrality* [27] counts the number of the node's neighbours: $D_v(G) = |\mathcal{N}_G(v)|$,
- *Closeness centrality* [3] is the inverse of the sum of distances to all of the other nodes:

$$C_v(G) = \frac{1}{\sum_{u \in V \setminus \{v\}} dist_G(u, v)},$$

- *Betweenness centrality* [14] measures the fraction of the shortest paths that traverse a node:

$$B_v(G) = \sum_{u, w \in V \setminus \{v\}} \frac{|\{p \in \Pi_G(u, w) : v \in p\}|}{|\Pi_G(u, w)|},$$

- *Eigeinvector centrality* [5] relies on the assumption that the centrality of a node is proportional to the sum of the centralitites of its neighbours, i.e., it is a solution to the recursive formula: $E_v(G) = \frac{1}{\lambda} \sum_{u \in \mathcal{N}_G(v)} E_u(G)$.

For a centrality measure, F, we consider a ranking of a node, $v \in V$, that results from F and denote it by $r^F_v(G)$.

Network models: To better understand the properties of networks, number of probabilistic models of networks have been proposed in the literature [2, 11, 32]. Such a network model can be understood as a discrete probability distribution, $\mathcal{G}$, over a space of all possible graphs $\mathbb{G}^V$. Hence, in the remainder of the paper, by G we will usually denote a *random variable* that is a graph drawn from $\mathcal{G}$. Specific graphs, e.g., realisations of this random variable, will be denoted by G with lower indices, e.g., G_0.

For two graph distributions $\mathcal{G}$ and $\mathcal{G}'$ and two constants $x, y \geq 0$ such that $x + y = 1$ we can take a convex combination of two graph distributions $x\mathcal{G} + y\mathcal{G}'$, such that for every $G_0 \in \mathbb{G}^V$ we have

$$\mathbb{P}_{x\mathcal{G}+y\mathcal{G}'}(G = G_0) = x\mathbb{P}_{\mathcal{G}}(G = G_0) + y\mathbb{P}_{\mathcal{G}'}(G = G_0).$$

3 PROBLEM OF MANIPULABILITY

We study the difficulty with which a node—the *evader*—can affect its centrality in the network by manipulating the network structure. The goal of the evader can be either hiding, i.e., decreasing its centrality ranking, or exposing—increasing its ranking. For the clarity of the presentation throughout the paper we will focus on the case of hiding, however all of our results can be easily extended to the case of exposing.

To study the problem of manipulability, we need additional terminology: *action*, *impact set* and *measure of manipulability*. By an *action* we will understand a possible change to a graph. Such an action can be either adding a new edge to a graph or removing an existing one. In both cases, for a graph $G = (V, E)$ we will denote an action as a 2-element subset of V (just like we denote an edge). We will denote by $a(G)$ the result of *performing* action a, i.e.,

$$a(V, E) = \begin{cases} (V, E \setminus a) & \text{if } a \in E, \\ (V, E \cup a) & \text{if } a \notin E. \end{cases}$$

Since performing actions is commutative, for a set of actions $S = \{a_1, \dots, a_k\}$ we will denote by $S(G)$ the result of performing all of the actions in S, i.e., $S(G) = a_1(\dots(a_k(G))\dots)$.

In the literature different sets of possible actions are studied, e.g., Crescenzi *et al.* [9] considered adding new edges to a specific node and Waniek *et al.* [31] studied removing edges of a node and adding edges between its neighbours. In this paper, we do not constrain ourselves to one particular set of possible actions. Instead, we allow for any arbitrary rule that characterise which actions are permitted in which graph. This approach makes our findings relevant to wider range of applications and allows for numerical comparison between evaders with different possibilities.

To this end, we define an *action function*, $\mathcal{A} : \mathbb{G}^V \to 2^{\{a \subseteq V : |a|=2\}}$, that for a given graph returns the set of possible actions that the evader can use. Here are a few examples of such action functions:

- $\mathcal{A}_0(G) = \{a \subseteq V : |a| = 2\}$ indicates that each edge can be removed or added in every graph;
- $\mathcal{A}_1(G) = \{a \in E[G] : v \in a\}$ indicates that in each graph existing edges of the evader can be removed;
- $\mathcal{A}_2(G) = \{a \subseteq \mathcal{N}_v(G) : |a| = 2 \land a \notin E[G]\}$ indicates that edges can be added between the neighbours of node v.

For an arbitrary action function $\mathcal{A}$ and action a, by $\mathcal{A} - a$ we will denote the action function given by $\mathcal{A} - a(G) = \mathcal{A}(G) \setminus \{a\}$.

Now, for a specified network $G_0 \in \mathbb{G}^V$, evader $v \in V$, centrality measure $F : \mathbb{G}^V \to \mathbb{R}^V_+$, and set of actions A, we will be interested in sets of actions that hide the evader, i.e., actions that decrease the ranking of v according to F.[1] We formalize this problem in the following definition.

Definition 3.1. (Impact Set) For a given graph G_0, its node v, centrality measure F and set of actions A, the *impact set* is the collection of all subsets of A that result in a decrease in the ranking of v when performed. Formally,

$$I^F_{G_0,v}(A) = \{S \subseteq A : r^F_v(G_0) > r^F_v(S(G_0))\}.$$

[1] Similarly, we could also consider exposing oneself in the network instead of hiding, i.e., the sets of action that increase the ranking of v. Another possibility is to consider decreasing (or increasing) the ranking of v by some fixed threshold.

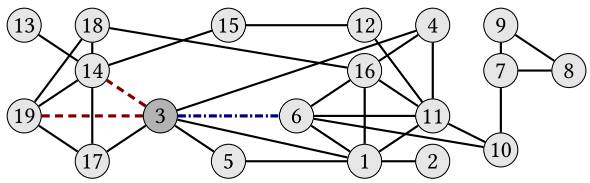

Figure 1: World Trade Center 9/11 terrorist network [18]. Node 3 (which represents the terrorist Nawaf Alhazmi) has to remove two of its edges (e.g., $\{3, 14\}$ and $\{3, 19\}$) to decrease its ranking according to Degree centrality. On the other hand, by removing just one edge (e.g., $\{3, 6\}$) it can decrease its ranking according to Closeness centrality.

Example 3.2. Let us consider the World Trade Center 9/11 terrorist network [18], which is illustrated in Figure 1, and let us denote it by G_T. Here, Degree centrality of node 3 (which represents the terrorist Nawaf Alhazmi) is equal to 7 and its Closeness centrality is equal to 1/35. Therefore, it is ranked as first according to both centrality measures. Let us consider the set of possible actions that allows for removing edges of node 3, i.e., $A_1 = \{a \in E[G_T] : 3 \in a\}$.

For Degree centrality, node 3 must remove at least 2 of its edges in order to decrease its centrality below Degree centrality of the next node in the ranking—node 11 (which represents Marwan Al-Shehhi). Hence, the impact set for graph G_T, node 3, Degree centrality, and set of actions A_1 will be equal to the collection of all sets of at least two edges incident with node 3 or, in other words, subsets of A_1 with at least two elements, i.e., $I^D_{G_T,3}(A_1) = \{S \subseteq A_1 : |S| \geq 2\}$.

For Closeness centrality, it turns out that removing only one edge from $\{3, 4\}, \{3, 5\}, \{3, 6\}$, and $\{3, 14\}$ is enough to decrease the ranking of node 3, however after removing edges $\{3, 14\}$ and $\{3, 19\}$ it is still ranked as first. Hence, the impact set consists of all possible subsets of allowed actions except for empty set, single edge $\{3, 14\}$, edge $\{3, 19\}$, and both edges combined, i.e., $I^C_{G_T,3}(A_1) = \big\{S \subseteq A_1\big\} \setminus \big\{\emptyset, \{\{3, 14\}\}, \{\{3, 19\}\}, \{\{3, 14\}, \{3, 19\}\}\big\}$.

Finally, to assess the easiness with which the evader can hide herself, we introduce the concept of a *measure of manipulability.*

Definition 3.3. A *measure of manipulability* is a function, M, that for a every graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measure F, and action function $\mathcal{A}$ returns a real value from the interval $[0, 1]$.

The value returned by such a measure represents the manipulability of the centrality measure given a network model with one node designated as an evader and an action function. The greater the value, the easier it is for the evader to hide herself in this network.

4 AXIOMS FOR A MEASURE OF MANIPULABILITY

Definition 3.3 of a measure of manipulability is very broad. To focus on more desirable measures of manipulability, we propose properties, i.e., axioms, that a measure of manipulability should satisfy.

The first such an axiom, *Unmanipulability*, states that if it is certain that no combination of actions will hide an evader, then the manipulability is zero.

> ***Unmanipulability:*** *For every graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measure F, and action function $\mathcal{A}$, if $\mathbb{P}\big(I^F_{G,v}(\mathcal{A}(G)) = \emptyset\big) = 1$, then $M(\mathcal{G}, v, F, \mathcal{A}) = 0$.*

Conversely, the second axiom, *Full Manipulability*, states that in the case, when it is sure to hide by any nonempty set of possible actions, the manipulability is equal to 1.

> ***Full Manipulability:*** *For every graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measure F, and action function $\mathcal{A}$, if*
> $$\mathbb{P}\big(I^F_{G,v}(\mathcal{A}) = \{S \subseteq \mathcal{A}(G) : S \neq \emptyset\}\big) = 1,$$
> *then $M(\mathcal{G}, v, F, \mathcal{A}) = 1$.*

Now, consider a scenario in which one centrality measure and action function are dominated by the other, i.e., whenever certain set of actions hides the evader under one measure it hides the evader also under the other measure. Our next axiom, *Weak Dominance*, states that in such a case the manipulability of the 'dominant' centrality measure and action function is greater or equal to the manipulability of the 'dominated' ones.

> ***Weak Dominance:*** *For every graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measures F and F', and action functions $\mathcal{A}$ and $\mathcal{A}'$, if*
> $$\mathbb{P}\big(I^F_{G,v}(\mathcal{A}) \subseteq I^{F'}_{G,v}(\mathcal{A}')\big) = 1,$$
> *then $M(\mathcal{G}, v, F, \mathcal{A}) \leq M(\mathcal{G}, v, F', \mathcal{A}')$.*

Our next axiom, *Neutrality*, states that a measure of manipulability should not unreasonably prefer one graph, node or centrality measure over the other. Hence, if we exchange graph or node or centrality measure, but the sets of action that hides the evader are still exactly the same, the manipulability is the same as well.

> ***Neutrality:*** *For every node $v \in V$, bijections $f : V \to V$ and $g : \mathbb{G}^V \to \mathbb{G}^V$ centrality measures F and F', and action function $\mathcal{A}$ if*
> $$I^F_{G,v}(\mathcal{A}) = I^{F'}_{g(G),f(v)}(\mathcal{A}) \quad \text{for every } G \in \mathbb{G}^V$$
> *then $M(\mathcal{G}, v, F, \mathcal{A}) = M(\mathcal{G}', f(v), F', \mathcal{A})$, for every graph distributions $\mathcal{G}$ and $\mathcal{G}'$ on space $\mathbb{G}^V$ such that $\mathbb{P}_{\mathcal{G}}(G = g(G_0)) = \mathbb{P}_{\mathcal{G}'}(G = G_0)$ for every $G_0 \in \mathbb{G}^V$.*

The next axiom considers an action, e.g., a, that is *redundant* for hiding an evader, i.e., whenever some set of actions that includes a hides an evader, the same set without a or with another action instead of a still hides the evader. This property can be formalised in the following definition.

Definition 4.1. For graph $G_0 \in \mathbb{G}^V$, node $v \in V$, centrality measure F and set of actions A, we say that action $a \in A$ is *redundant*, if for every subset of actions $S \subseteq A$ that hides node v, i.e., $S \in I^F_{G,v}(A)$, the fact that $a \in S$ implies that either $S \setminus \{a\} \in I^F_{G,v}(A)$ or there exist another action $a' \in A \setminus S$ such that $S \setminus \{a\} \cup \{a'\} \in I^F_{G,v}(A)$.

Now, the axiom *Redundant Action* states that if a particular action is always redundant, then excluding it from action function does not affect the manipulability. This property is important if we think about strategic manipulation: For the evader that can purposefully choose its hiding strategy redundant actions are irrelevant.

> ***Redundant Action:*** *For every graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measures F, and action function $\mathcal{A}$, if there exist an action such that $\mathbb{P}(a \in \mathcal{A}(G)) = 1$ and $\mathbb{P}(a \text{ is redundant}) = 1$, then $M(\mathcal{G}, v, F, \mathcal{A}) = M(\mathcal{G}, v, F, \mathcal{A} - a)$.*

Our next axiom, *Linearity*, states that manipulability over the combination of two network models is a combination of manipulabilities over these network models.

> ***Linearity:*** *For every two graph distributions $\mathcal{G}$ and $\mathcal{G}'$ on space $\mathbb{G}^V$, node $v \in V$, centrality measure F, action function $\mathcal{A}$, and two constants $x, y > 0$ such that $x + y = 1$ it holds that*

$$M(x\mathcal{G} + y\mathcal{G}', v, F, \mathcal{A}) = xM(\mathcal{G}, v, F, \mathcal{A}) + yM(\mathcal{G}', v, F, \mathcal{A}).$$

Our first six axioms, i.e., Unmanipulability, Full Manipulability, Weak Dominance, Neutrality, Redundant Action, and Linearity characterise a family of measures of manipulability that depend solely on minimal number of actions required to hide the evader (see Theorem 5.4). Adding our last axiom, *Normalisation*, uniquely characterise Average Minimal Action Required (AMAR) measure of manipulability (see Theorem 5.7).

> ***Normalisation:*** *For every constant $k \in \{1, \dots, (|V|^2 - |V|)/2\}$, there exist graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measure F, and action function $\mathcal{A}$ such that $|\mathcal{A}(G_0)| \geq k$ for every $G_0 \in \mathbb{G}^V$,*
>
> $$\mathbb{P}(I^F_{G,v}(\mathcal{A}(G)) = \{S \subseteq \mathcal{A}(G) : |S| = k\}) = 1$$
>
> *and $M(\mathcal{G}^V, v, F, \mathcal{A}) = 1/k$.*

5 THE AMAR MEASURE OF MANIPULABILITY

In this section we introduce a measure of manipulability that satisfies all of our axioms. We will call it *Average Minimal Actions Required (AMAR)* and define it as the average inverse of the minimal number of actions required to hide the evader. The cases in which it is impossible to hide the evader are counted as 0.

To formally define Average Minimal Actions Required let us begin with the definition of Minimal Actions Required (MAR), which for each graph, a node, a centrality measure and a set of actions returns a value between 0 and 1 that represent the easiness of manipulation in this particular setting.

Definition 5.1. Let *Minimal Actions Required (MAR)* be defined by the formula:

$$MAR(G, v, F, A) = \begin{cases} 0 & \text{if } I^F_{G,v}(A) = \emptyset, \\ \frac{1}{\min_{S \in I^F_{G,v}(A} |S|} & \text{otherwise.} \end{cases} \quad (1)$$

Example 5.2. Recall example 3.2 with 9/11 terrorist network, G_T, depicted on Figure 1 and set of possible actions that allowed for removing edges incident with node 3, i.e., $A_1 = \{a \in E[G_T] : v \in a\}$. In order to hide itself, node 3 must remove at least 2 of its edges, to decrease its ranking according to Degree centrality, i.e., $I^D_{G_T,3} = \{S \subseteq A_1 : |S| \geq 2\}$. Hence, the MAR measure of manipulability is equal to $MAR(G_T, 3, D, A_1) = 1/2$.

For Closeness centrality, it is possible that node 3 hides itself by removing just one edge, e.g., $\{3, 6\}$. Hence, $MAR(G_T, 3, C, A_1) = 1$.

Building upon Minimal Actions Required function, we can define Average Minimal Actions Required (AMAR) measure of mainpulability.

Definition 5.3. Let *Average Minimal Actions Required (AMAR)* be a measure of manipulability defined by the formula:

$$AMAR(\mathcal{G}, v, F, \mathcal{A}) = \mathbb{E}_{\mathcal{G}}\big(MAR(G, v, F, \mathcal{A}(G))\big). \quad (2)$$

In Theorem 5.7 we will prove that AMAR measrue of manipulability is the unique measure of manipulability that satisfies all of our axioms. However, we begin with the analysis of our first six axioms, i.e., Unmanipulability, Full Manipulability, Weak Dominance, Neutrality, Redundant Action, and Linearity, and prove that the measure of manipulability that satisfies them depends solely on the minimal number of required actions in each graph, i.e., MAR.

THEOREM 5.4. *A measure of manipulability, M, satisfies Unmanipulability, Full Manipulability, Weak Dominance, Neutrality, Redundant Action, and Linearity if and only if there exists a nondecreasing function, $f : [0, 1] \to [0, 1]$, such that $f(0) = 0$, $f(1) = 1$, and*

$$M(\mathcal{G}, v, F, \mathcal{A}) = \mathbb{E}_{\mathcal{G}}\big(f\big(MAR(G, v, F, \mathcal{A}(G))\big)\big) \quad (3)$$

for every graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measure F and action set $\mathcal{A}$.

PROOF. It is easy to check that the measure of manipulability defined by the equation (3) for a nondecreasing function, $f : [0, 1] \to [0, 1]$, such that $f(0) = 0$, $f(1) = 1$ satisfies Unmanipulability, Full Manipulability, Weak Dominance, Neutrality, Redundant Action, and Linearity. Therefore, in the proof we will focus on showing that the measure of manipulability that satisfies these axioms imply that there exists a nondecreasing function, $f : [0, 1] \to [0, 1]$, such that $f(0) = 0$, $f(1) = 1$ such that $M(\mathcal{G}, v, F, \mathcal{A}) = \mathbb{E}_{\mathcal{G}}\big(f\big(MAR(G, v, F, \mathcal{A}(G))\big)\big)$.

To this end, we will need additional notation: For every graph $G_0 \in \mathbb{G}^V$, a *single graph distribution*, denoted by δ_{G_0}, is a graph distribution fixed at graph G_0, i.e., such that $\mathbb{P}_{\delta_{G_0}}(G = G_0) = 1$.

Let us focus on single graph distributions in the following lemma in which we also assume that all sets with at least k actions hides the evader.

LEMMA 5.5. *If a measure of manipulability, M, satisfies Weak Dominance, Neutrality, and Redundant Action, then there exist a function $f : \mathbb{N} \to [0, 1]$ such that*

$$M(\delta_{G_0}, v, F, \mathcal{A}) = f(k)$$

for every graph $G_0 \in \mathbb{G}^V$, node $v \in V$, centrality measure F and action function $\mathcal{A}$ such that

$$I^F_{G_0,v}(\mathcal{A}(G_0)) = \{S \subseteq \mathcal{A}(G_0) : |S| \geq k\} \neq \emptyset. \quad (4)$$

PROOF. Fix $k \in \mathbb{N}$. We will prove this lemma by showing that there exists one specific $G^*_0 \in \mathbb{G}^V$, $v^* \in V$, F^* and $\mathcal{A}^*$ such that for each $G_0 \in \mathbb{G}^V$, $v \in V$, F and $\mathcal{A}$ that satisfies condition (4) it holds that

$$M(\delta_{G_0}, v, F, \mathcal{A}) = M(\delta_{G_0^*}, v^*, F^*, \mathcal{A}^*). \quad (5)$$

But first, for every graph $G_0 \in \mathbb{G}^V$, node $v \in V$ and $k \in \mathbb{N}$ let us construct an auxiliary centrality measure $F^{G_0,v,k}$ that is constant for every node except v for which it changes if there are more than k changes made to graph G_0. Formally, for every graph $G \in \mathbb{G}^V$ and node $u \in V$ let

$$F_u^{G_0,v,k}(G) = \begin{cases} 1 & \text{if } u = v \text{ and } \left|E[G] \ominus E[G_0]\right| \geq k, \\ 2 & \text{otherwise,} \end{cases}$$

where $\ominus$ is a symmetric difference between sets. In this way, for an arbitrary set of actions $S \subseteq \{a \subseteq V : |a| = 2\}$

$$r_v^{F^{G_0,v,k}}(S(G_0)) = \begin{cases} 1 & \text{if } |S| < k, \\ n & \text{if } |S| \geq k. \end{cases}$$

As a result, every subset of k or more actions (and nothing more) hides v in G_0 according to $F^{G_0,v,k}$, i.e., for every set of actions A, impact set $I_{G_0,v}^{F^{G_0,v,k}}(A)$ consists of all at least k element subsets of A.

Now, let us denote an action function of all possible actions as $\mathcal{A}^* = \{a \subseteq V : |a| = 2\}$. We will prove that for every F and $\mathcal{A}$ that satisfies condition (4) the manipulability of v in G_0 is equal the manipulability of v in G_0 for $F^{G_0,v,k}$ and $\mathcal{A}^*$, i.e.,

$$M(\delta_{G_0}, v, F, \mathcal{A}) = M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}^*). \quad (6)$$

To this end, we show that $M(\delta_{G_0}, v, F, \mathcal{A}) = M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A})$ and then that $M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}) = M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}^*)$.

As stated, for every set of actions A, impact set $I_{G_0,v}^{F^{G_0,v,k}}(A)$ consists of all at least k element subsets of A. Therefore, for action set $\mathcal{A}(G_0)$ we have $I_{G_0,v}^{F^{G_0,v,k}}\mathcal{A}(G_0) = \{S \subseteq \mathcal{A}(G_0) : |S| \geq k\}$. Hence, from the fact that F and $\mathcal{A}$ satisfy condition 4 we get $I_{G_0,v}^{F^{G_0,v,k}}\mathcal{A}(G_0) = I_{G_0,v}^{F}\mathcal{A}(G_0)$. Thus, $\mathbb{P}_{\delta_{G_0}}(I_{G_0,v}^{F^{G_0,v,k}}\mathcal{A}(G_0) \subseteq I_{G_0,v}^{F}\mathcal{A}(G_0))$ as well as $\mathbb{P}_{\delta_{G_0}}(I_{G_0,v}^{F^{G_0,v,k}}\mathcal{A}(G_0) \supseteq I_{G_0,v}^{F}\mathcal{A}(G_0))$. Therefore, from Weak Dominance we get

$$M(\delta_{G_0}, v, F, \mathcal{A}) = M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}). \quad (7)$$

Now, let us prove that the manipulability is still the same if we exchange $\mathcal{A}$ for the action function of all possible actions $\mathcal{A}^*$, i.e.,

$$M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}) = M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}^*). \quad (8)$$

We will prove this equation by induction on the difference between number of actions in $\mathcal{A}^*(G_0)$ and $\mathcal{A}(G_0)$. If $|\mathcal{A}^*(G_0)| - |\mathcal{A}(G_0)| = 0$, then $\mathcal{A}(G_0)$ must contain all possible actions as well, i.e., $\mathcal{A}(G_0) = \{a \subseteq V : |a| = 2\} = \mathcal{A}^*(G_0)$. Hence, equation (8) holds.

Now, assume that for some $m \in \mathbb{N}$ equation (8) holds for all action functions $\mathcal{A}$ such that $|\mathcal{A}^*(G_0)| - |\mathcal{A}(G_0)| = m$. Let us take an arbitrary action function $\mathcal{A}$ such that $|\mathcal{A}^*(G_0)| - |\mathcal{A}(G_0)| = m + 1$. Consider an action a that belongs to set $\mathcal{A}^*(G_0)$, but does not belong to $\mathcal{A}(G_0)$, i.e., $a \in \mathcal{A}^*(G_0) \setminus \mathcal{A}(G_0)$. Also, let us denote action function $\mathcal{A}'$ as a function $\mathcal{A}$ with action a added, i.e., $\mathcal{A}'(G) = \mathcal{A}(G) \cup \{a\}$ for every $G \in \mathbb{G}^V$. With this additional action $|\mathcal{A}^*(G_0)| - |\mathcal{A}'(G_0)| = m$, hence from our inductive assumption

$$M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}') = M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}^*). \quad (9)$$

On the other hand, consider a set of actions that hides node v, i.e., $S \in I_{G_0,v}^{F}(\mathcal{A}'(G_0))$, such that $a \in S$. If $|S| > k$, then from condition (4) the same set without action a still hides node v, i.e., $S \setminus \{a\} \in I_{G_0,v}^{F}(\mathcal{A}'(G_0))$. If $|S| = k$, then S does not contain all of the actions in $\mathcal{A}'(G_0)$. Hence, there exists an action $a' \in \mathcal{A}'(G_0)$, $a' \notin S$ and from condition (4) set of actions $S \setminus \{a\} \cup \{a'\}$ hides node v as well, i.e., $S \setminus \{a\} \cup \{a'\} \in I_{G_0,v}^{F}(\mathcal{A}'(G_0))$. Hence, action a is redundant. Since it is sure that $G = G_0$ it holds that $\mathbb{P}_{\delta_{G_0}}(a \text{ is redundant}) = 1$. Therefore, from Redundant Action

$$M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}) = M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}'). \quad (10)$$

By combining equations (9) and (10) we obtain equation (8) and from induction we get that it holds for all $\mathcal{A}$ such that $|\mathcal{A}(G_0)| \geq k$. Finally, from equation (7) and equation (8) we get

$$M(\delta_{G_0}, v, F, \mathcal{A}) = M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}^*)$$

for every centrality measure F and action function $\mathcal{A}$ such that $I_{G_0,v}^{F}(\mathcal{A}(G_0)) = \{S \subseteq \mathcal{A}(G_0) : |S| \geq k\} \neq \emptyset$.

Now, it remains to prove that for some one specific G^*, v^* and F^* it holds that $M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}^*) = M(\delta_{G^*}, v^*, F^*, \mathcal{A}^*)$. Let $G^* = (V, \emptyset)$, v^* be a specific node in V and $F^* = F^{G*,v^*,k}$. Consider bijection $g : \mathbb{G}^V \rightarrow \mathbb{G}^V$ such that

$$g(V, E) = (V, E \ominus E[G_0])$$

where $\ominus$ is a symmetric difference. Observe that $g(G_0) = G^*$ and

$$E[g(G)] \ominus E[G_0] = E[G] \ominus E[G_0] \ominus E[G_0] = E[G] \ominus E[G^*] \quad (11)$$

Moreover, let us consider some bijection $h : V \rightarrow V$ such that $h(v) = v^*$. Then, from equation (11) we obtain that $F_u^{G_0,v,k}(g(G)) = F_{f(u)}^{G^*,v^*,k}(G)$ for every $G \in \mathbb{G}^V$ and $u \in V$. Hence, for any action function $\mathcal{A}$ we get that $I_{G,v}^{F^{G_0,v,k}}(\mathcal{A}(G)) = I_{G,h(v)}^{F^{G^*,h(v),k}}(\mathcal{A}(G))$ for every $G \in \mathbb{G}^V$. Thus, the thesis follows from Neutrality. □

In Lemma 5.5 we focused on single graph distributions and assumed that all sets of k or more actions hide the evader. In the following lemma we relax the latter assumption.

Lemma 5.6. *If a measure of manipulability, M, satisfies Unmanipulability, Full Manipulability, Weak Dominance, Neutrality, and Redundant Action, then there exists nondecreasing function $f : [0,1] \rightarrow [0,1]$ such that $f(0) = 0$, $f(1) = 1$ and*

$$M(\delta_{G_0}, v, F, \mathcal{A}) = f\left(MAR(G_0, v, F, \mathcal{A}(G_0))\right) \quad (12)$$

for every $G_0 \in \mathbb{G}^V$, $v \in V$, centrality F and action function $\mathcal{A}$.

Proof. Let us assume that measure of manipulability M satisfies Unmanipulability, Full Manipulability, Weak Dominance, Neutrality and Redundant Action. Fix $G_0 \in \mathbb{G}^V$, node $v \in V$, centrality measure F and action function $\mathcal{A}$. From Lemma 5.5 we know that there exists a function, $f : [0,1] \rightarrow [0,1]$, such that for every single graph distribution, if all sets of k or more action hide the evader, then the manipulability is equal to $f(1/k)$. Now, we will prove that for the same function f we have

$$M(\delta_{G_0}, v, F, \mathcal{A}) = \begin{cases} 0 & \text{if } I_{G_0,v}^{F}(\mathcal{A}(G_0)) = \emptyset, \\ f\left(\min_{S \in I_{G_0,v}^{F}(\mathcal{A}(G_0))} |S|\right) & \text{otherwise.} \end{cases}$$

If $I^F_{G_0,v}(\mathcal{A}(G_0))=\emptyset$, from Unmanipulability $M(\delta_{G_0}, v, F, \mathcal{A})=0$.

Hence, let us focus on a case in which $I^F_{G_0,v}(\mathcal{A}(G_0)) \neq \emptyset$. By k let us denote the size of minimal set of actions that hides the evader, i.e., let $k = \min_{S \in I^F_{G_0,v}(\mathcal{A}(G_0))} |S|$. Also, by T let us denote one of these minimal sets, i.e., let $T \in I^F_{G_0,v}(\mathcal{A}(G_0))$ such that $|T| = k$. Let us denote also the action function $\mathcal{T}$ that for every graph returns set T. Observe that from Weak Dominance we get that $M(\delta_{G_0}, v, F, \mathcal{A}) \geq M(\delta_{G_0}, v, F, \mathcal{T})$. Furthermore, since $I^F_{G_0,v}(\mathcal{T}(G_0)) = \{T\}$ and $|T| = k$, from Lemma 5.5 we get that

$$M(\delta_{G_0}, v, F, \mathcal{A}) \geq f(k). \tag{13}$$

On the other hand, let us consider once again centrality measure $F^{G_0,v,k}$ and action function $\mathcal{A}^*$ from the proof of Lemma 5.5. Let $\mathcal{A}^*$ be an action function of all possible actions, i.e., $\mathcal{A}^*(G) = \{a \subseteq V : |a| = 2\}$, and let $F^{G_0,v,k}$ be a centrality measure depending on number of changes to graph G_0, i.e.,

$$F^{G_0,v,k}_u(G) = \begin{cases} 1 & \text{if } u = v \text{ and } \left|E[G] \ominus E[G_0]\right| \geq k, \\ 2 & \text{otherwise.} \end{cases}$$

In such a case sets of actions hiding v in G_0 under $F^{G_0,v,k}$ are the sets of actions that consists of k or more elements. This means that $I^{F^{G_0,v,k}}_{G_0,v}(\mathcal{A}^*(G_0)) = \{S \subseteq \mathcal{A}*(G_0) : |S| \geq k\}$. Hence, from Lemma 5.5 we get that

$$M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}^*) = f(k). \tag{14}$$

Now, observe that $\mathcal{A}(G_0) \subseteq \mathcal{A}^*(G_0)$. Furthermore, for every $S \in I^F_{G_0,v}(\mathcal{A}(G_0))$ we know that $|S| \geq k$, hence $S \in I^{F^{G_0,v,k}}_{G_0,v}(\mathcal{A}^*(G_0))$. Thus, $I^F_{G_0,v}(\mathcal{A}(G_0)) \subseteq I^{F^{G_0,v,k}}_{G_0,v}(\mathcal{A}^*(G_0))$ and from Weak Dominance we get $M(\delta_{G_0}, v, F, \mathcal{A}) \leq M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}^*)$. With equation (14) this yields

$$M(\delta_{G_0}, v, F, \mathcal{A}) \leq f(k). \tag{15}$$

Finally, combining inequalities (13) and (15) we get

$$M(\delta_{G_0}, v, F, \mathcal{A}) = f(k).$$

It remains to prove that $f(1) = 1$ and f is a nondecreasing function. To see why $f(1) = 1$, observe that if $k = 1$ then from equation (14) we have $M(\delta_{G_0}, v, F^{G_0,v,1}, \mathcal{A}^*) = f(1)$. On the other hand, for $F^{G_0,v,1}$ and $\mathcal{A}^*$ every nonempty set of actions hides node v, hence from Full Manipulability $f(1) = 1$.

To see why f is nondecreasing consider centrality measures $F^{G_0,v,k}$ and $F^{G_0,v,k+1}$ and action function $\mathcal{A}^*$. Observe that

$$I^{F^{G_0,v,k+1}}_{G_0,v}(\mathcal{A}^*(G_0)) = \{S \subseteq \mathcal{A}^*(G_0) : |S| \geq k+1\} \subset \{S \subseteq \mathcal{A}^*(G_0) : |S| \geq k\} = I^{F^{G_0,v,k}}_{G_0,v}(\mathcal{A}^*(G_0)).$$

Hence, from Weak Dominance it holds that

$$M(\delta_{G_0}, v, F^{G_0,v,k+1}, \mathcal{A}^*) \leq M(\delta_{G_0}, v, F^{G_0,v,k}, \mathcal{A}^*).$$

Combining it with equation (14) yields $f(k+1) \leq f(k)$ for arbitrary $k \in \mathbb{N}$. Therefore, f is nonincreasing. □

In Lemmas 5.5 and 5.6 we focused on single graph distributions. In the remainder of the proof, we will generalise these results to arbitrary distributions. To this end, we use induction on the number of different graphs that have non-zero probability of being picked.

Let us consider arbitrary graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measure F and action function $\mathcal{A}$. If there exist only one graph $G_1 \in \mathbb{G}^V$ such that $\mathbb{P}_{\mathcal{G}}(G \in \{G_1\}) = 1$, then $\mathcal{G} = \delta_{G_1}$ and the thesis follows from Lemma 5.6.

Now, let us assume that for some $k \in \mathbb{N}$, the thesis holds for every graph distribution $\mathcal{G}$ for which there exist k graphs $G_1, \ldots, G_k \in \mathbb{G}^V$ such that $\mathbb{P}_{\mathcal{G}}(G \in \{G_1, \ldots, G_k\}) = 1$. Let us consider such graph distribution $\mathcal{G}$ for which there exist $k+1$ graphs $G_0, \ldots, G_k \in \mathbb{G}^V$ such that $\mathbb{P}_{\mathcal{G}}(G \in \{G_0, \ldots, G_k\}) = 1$. We denote the probability that graph G_0 is drawn from distribution $\mathcal{G}$ by $p_0 = \mathbb{P}_{\mathcal{G}}(G = G_0)$. Next, consider graph distribution $\mathcal{G}'$ which is distribution $\mathcal{G}$ without graph G_0. Formally, for every $i \in \{1, \ldots, k\}$ it holds that

$$\mathbb{P}_{\mathcal{G}'}(G = G_i) = \frac{\mathbb{P}_{\mathcal{G}}(G = G_i)}{1 - p_0}.$$

Observe that $\mathcal{G}$ is a convex combination of $\mathcal{G}'$ and δ_{G_0}, i.e., $\mathcal{G} = (1 - p_0)\mathcal{G}' + p_0\delta_{G_0}$. Hence, from Linearity

$$M(\mathcal{G}, v, F, \mathcal{A}) = (1 - p_0)M(\mathcal{G}', v, F, \mathcal{A}) + p_0 M(\delta_{G_0}, v, F, \mathcal{A}). \tag{16}$$

Observe that $\mathbb{P}_{\mathcal{G}'}(G \in \{G_1, \ldots, G_k\}) = 1$. Hence, from inductive assumption we obtain

$$M(\mathcal{G}', v, F, \mathcal{A}) = \mathbb{E}_{\mathcal{G}'}\left(f\left(MAR(G, v, F, \mathcal{A}(G))\right)\right). \tag{17}$$

On the other hand, from Lemma 5.6 we get

$$M(\delta_{G_0}, v, F, \mathcal{A}) = \mathbb{E}_{\delta_{G_0}}\left(f\left(MAR(G, v, F, \mathcal{A}(G))\right)\right). \tag{18}$$

Combining equations (16), (17) and (18) yields $M(\mathcal{G}, v, F, \mathcal{A}) = \mathbb{E}_{\mathcal{G}}\left(f\left(MAR(G, v, F, \mathcal{A})\right)\right)$. Thus, the thesis follows from induction and the fact that for each set of nodes V space $\mathbb{G}^V$ is finite. □

Theorem 5.7. *If a measure of manipulability, M, satisfies Unmanipulability, Full Manipulability, Weak Dominance, Neutrality, Redundant Action, Linearity and Normalisation, then for every graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measure F and action function $\mathcal{A}$ it holds that $M(\mathcal{G}, v, F, \mathcal{A}) = AMAR(\mathcal{G}, v, F, \mathcal{A})$.*

Proof. It is easy to check that AMAR satisfies Normalisation and it satisfies the remaining axioms from Theorem 5.4.

From Normalisation we know that for every $k \in \{1, \ldots, (|V|^2 - |V|)/2\}$ there exist graph distribution $\mathcal{G}$ on space $\mathbb{G}^V$, node $v \in V$, centrality measure F and action function $\mathcal{A}$ such that $|\mathcal{A}(G_0)| \geq k$ for every $G_0 \in \mathbb{G}^V$, $\mathbb{P}(I^F_{G,v}(\mathcal{A}(G)) = \{S \subseteq \mathcal{A}(G) : |S| = k\}) = 1$ and $M(\mathcal{G}, v, F, \mathcal{A}) = 1/k$. From Theorem 5.4 we know that

$$M(\mathcal{G}, v, F, \mathcal{A}) = \mathbb{E}(f(MAR(G, v, F, \mathcal{A}))) = f(1/k).$$

Hence, for all $k \in \{1, \ldots, (|V|^2 - |V|)/2\}$ we have that $f(1/k) = 1/k$. Since these are all possible positive values of MAR, this concludes the thesis. □

We note again that all of the technical results hold also if we consider the problem of exposing oneself in the network instead of hiding for identically defined AMAR measure.

6 EMPIRICAL ANALYSIS

Having developed a measure of manipulability, we will now use it to quantify the manipulability of various centrality measures and network models. In our experiments, we consider networks generated using six network models:

- *Random graphs*, generated using the Erdős-Rényi model [11]. In our experiments we set the expected average degree in a network to be 4.
- *Small-world* networks, generated using the Watts-Strogatz model [32]. In our experiments we set the expected average degree to be 4 and the rewiring probability to be $\frac{1}{4}$.
- *Preferential attachment* networks, generated using the Barabási-Albert model [2]. In our experiments we add 2 edges with each new node, and we set the size of the initial clique to 2.
- *Random trees* generated using Prüfer sequences [24]. In our experiments we use sequences where each element is chosen uniformly at random from set $1, \ldots, n$.
- *Scale-free* networks generated using the configuration model [20]. In our experiments we assume the minimal degree to be 2, the maximal degree to be 6, and the configuration model parameter to be 3.
- *Cellular networks* are meant to reflect the structure of the covert organizations [29]. In our experiments we set the mean cell size to 6, cell density to 0.9, the density of connections between the leaders to 0.2, and the triad closure probability to 0.18.

For a given network, the evader is chosen as the node with the highest average ranking according to all four considered centrality measures. We then consider one of the following sets of action available to the evader:

- $\mathcal{A}_1$: *All changes*—the evader is allowed to perform any change in the network, i.e., $\mathcal{A}_1(G) = \{a \subseteq V : |a| = 2\}$;
- $\mathcal{A}_2$: *Remove neighbors*—the evader v is only allowed to remove edges between herself and her neighbors, i.e., $\mathcal{A}_2(G) = \{a \in E[G] : v \in a\}$;
- $\mathcal{A}_3$: *Add between neighbors*—the evader v is only allowed to add edges between her neighbors, i.e., $\mathcal{A}_3(G) = \{a \subseteq \mathcal{N}_G(v) : |a| = 2 \wedge a \notin E[G]\}$;
- $\mathcal{A}_4$: *Local changes*—the evader v is allowed to remove edges between herself and her neighbors, as well as to add and remove edges between her neighbors, i.e., $\mathcal{A}_4(G) = \{a \in E[G] : v \in a\} \cup \{a \subseteq \mathcal{N}_G(v) : |a| = 2\}$.

In our experiments, we generate networks where the number of nodes is between 8 and 50 for the *Remove neighbors* actions function, and between 8 and 20 for all other action functions (These functions typically result in much larger sets, thus it is much more computationally demanding to process the many subsets of those actions. This is also the reason why it is not feasible to calculate AMAR for larger networks). For a given network, evader, and set of actions, we compute the value of the MAR measure as follows. First, we consider all possible actions and check whether performing any of them decreases the ranking of the evader according to any of the centrality measures. Now, if there is at least one centrality measure for which we did not manage to decrease the ranking of the evader, then we consider all subsets consisting of two actions; if these also do not decrease the evader's ranking according to that centrality measure, then we consider all subsets consisting of three actions, and so on and so forth. We continue this process until eventually, for each centrality measure, we find a set of actions that decrease the ranking of the evader, or until we exhaust all subsets of five actions (due to implementation issues, if there exists no subset of five actions that decreases the evader's ranking according to a particular centrality measure, then MAR is assumed to return a value of zero for that centrality measure). We repeat the experiment 400 times for each combination of: network size, network model, and set of actions.

The results of our experiments are presented in Figure 2. In the majority of the scenarios, Degree centrality is significantly harder to manipulate than other centrality measures. The possible reason behind this phenomenon is that for Degree centrality, in order to change the position in the ranking between two nodes, one has to add/remove the number of edges equal to the difference in the degrees of these nodes. For other centrality measures it often suffices to add/remove only the most impactful edges (see Example 3.2). In case of the action function that allows only for adding edges between the neighbours, $\mathcal{A}_3$, in some scenarios, Eigeinvector centrality turns out to be the hardest to manipulate. In these scenarios, in order to hide itself the evader usually has to increase the centrality of one of its neighbours, so that it becomes greater then the centrality of the evader. This is especially hard in case of Eigeinvector centrality, since centrality of a node depends on the centrality of its neighbours.

As can be seen, for the networks generated using the Erdős-Rényi, Watts-Strogatz and configuration models, as well as Prüfer trees and cellular networks, the value of manipulability remains at about the same level for all investigated network sizes. Furthermore, in most cases this value is between 0.8 and 1, indicating that centrality measures in these types of networks can be manipulated relatively easily (in fact, in the majority of our experiments, it takes only a single action to change the evader's ranking). In contrast, for the networks generated using the Barabási-Albert model, the value of the manipulability measure decreases with the number of nodes, indicating that it is more challenging for hubs to hide in larger scale free networks. Moreover, on average the value of manipulability in the Barabási-Albert networks of a given size is in most cases lower than in networks of the same size generated using the other five models, suggesting that the centrality measures are more difficult to manipulate in networks with scale-free properties.

7 CONCLUSIONS

Centrality measures are among the most widely-used tools for social network analysis. A growing body of work focuses on understanding the susceptibility of such measures to manipulation by individuals who strategically rewire the network to their advantage, in the hope of misleading the centrality analysis. In this paper, we formalized the problem of quantifying the manipulability of centrality measures. We proposed a set of intuitive and seemingly-desirable axioms for such a measure. Based on this, we defined a measure that is uniquely characterized by our axioms. Finally, using our measure, we evaluated the manipulability of various centrality measures under different network models.

Ideas for future work include developing axiomatization of the manipulability of other kinds of social network analysis tools,

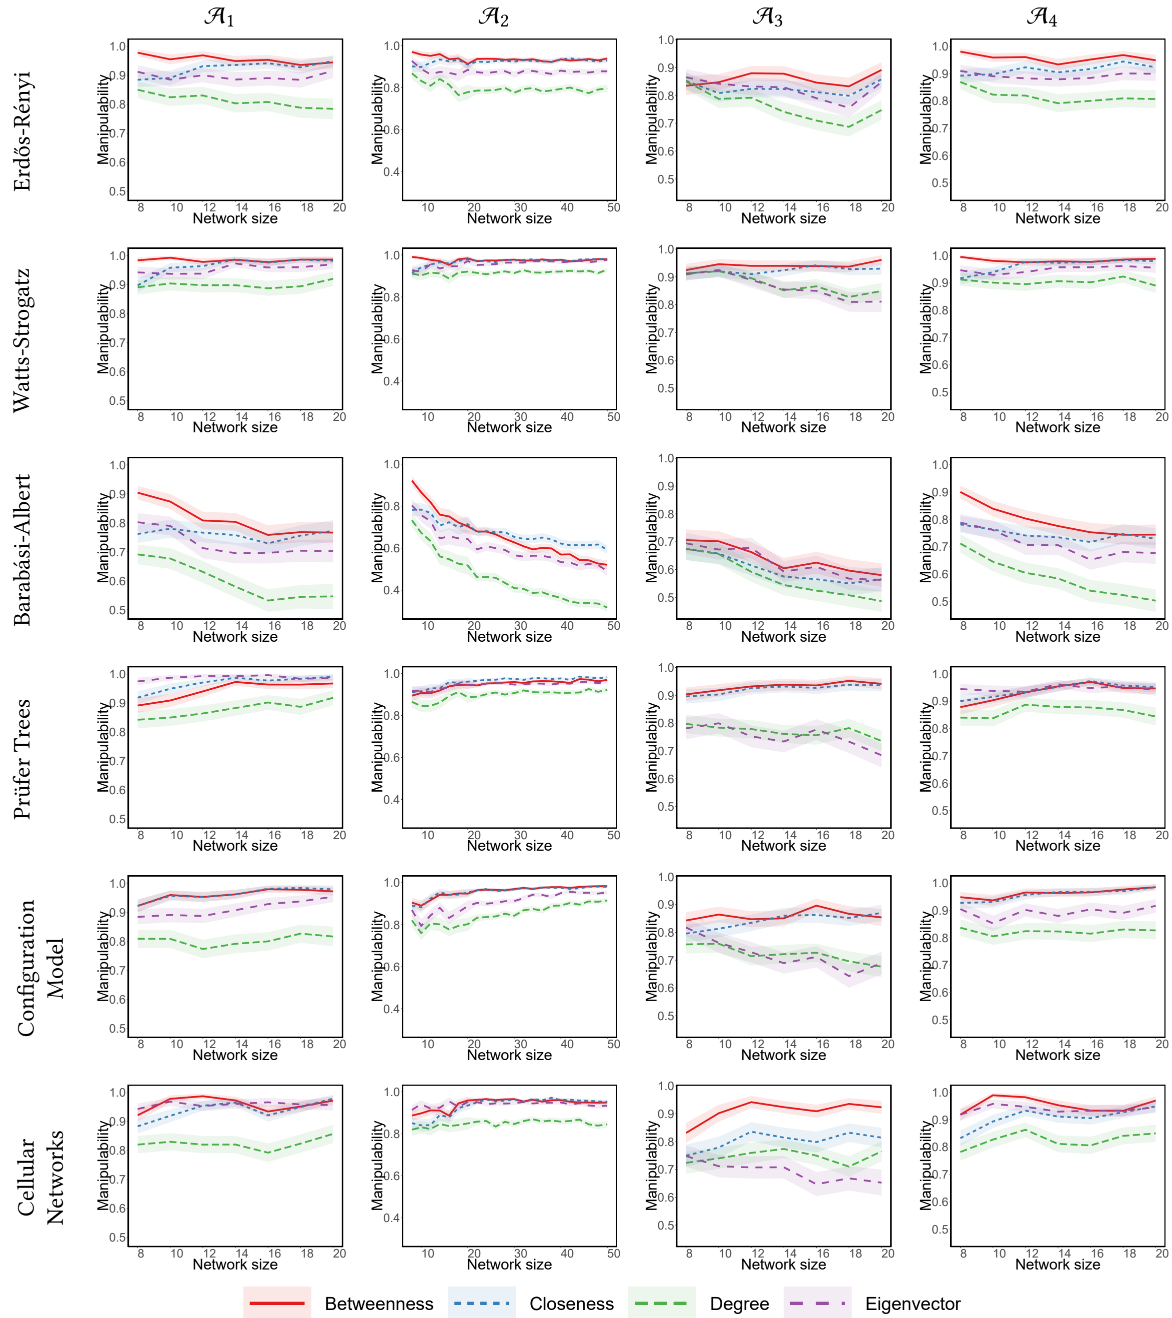

Figure 2: Results of our experiments with the AMAR measure of manipulability in randomly generated networks. Each row contains results for different network generation model, while each column contains results for different set of actions. Values represent AMAR manipulability measure estimated using 400 **networks. Colored areas represent** 95% **confidence intervals.**

e.g., link prediction algorithms and community detection algorithms. Another potential venue for extending our work is to investigate more sophisticated types of centrality measures (e.g., game-theoretic centrality measures) to determine whether they are less prone to manipulation than the four centrality measures considered in this paper.

ACKNOWLEDGMENTS

This work was supported by the Polish National Science Center under grant number 2016/23/B/ST6/03599. For the earlier version of this work Marcin Waniek was also supported by the Polish National Science Center under grant number 2015/17/N/ST6/03686.

REFERENCES

[1] Konstantin Avrachenkov and Nelly Litvak. 2006. The effect of new links on Google PageRank. *Stochastic Models* 22, 2 (2006), 319–331.

[2] Albert-László Barabási and Réka Albert. 1999. Emergence of scaling in random networks. *science* 286, 5439 (1999), 509–512.

[3] Alex Bavelas. 1950. Communication patterns in task-oriented groups. *The Journal of the Acoustical Society of America* 22, 6 (1950), 725–730.

[4] Elisabetta Bergamini, Pierluigi Crescenzi, Gianlorenzo D'angelo, Henning Meyerhenke, Lorenzo Severini, and Yllka Velaj. 2018. Improving the betweenness centrality of a node by adding links. *Journal of Experimental Algorithmics (JEA)* 23 (2018), 1–5.

[5] Phillip Bonacich. 1972. Factoring and weighting approaches to status scores and clique identification. *Journal of mathematical sociology* 2, 1 (1972), 113–120.

[6] Stephen P Borgatti, Kathleen M Carley, and David Krackhardt. 2006. On the robustness of centrality measures under conditions of imperfect data. *Social networks* 28, 2 (2006), 124–136.

[7] Yazan Boshmaf, Ildar Muslukhov, Konstantin Beznosov, and Matei Ripeanu. 2011. The socialbot network: when bots socialize for fame and money. In *Proceedings of the 27th annual computer security applications conference*. ACM, 93–102.

[8] Elizabeth Costenbader and Thomas W Valente. 2003. The stability of centrality measures when networks are sampled. *Social networks* 25, 4 (2003), 283–307.

[9] Pierluigi Crescenzi, Gianlorenzo D'angelo, Lorenzo Severini, and Yllka Velaj. 2016. Greedily improving our own closeness centrality in a network. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 11, 1 (2016), 9.

[10] Erik D Demaine and Morteza Zadimoghaddam. 2010. Minimizing the diameter of a network using shortcut edges. In *Scandinavian Workshop on Algorithm Theory*. Springer, 420–431.

[11] Paul Erdős and Alfréd Rényi. 1959. On random graphs I. *Publ. Math. Debrecen* 6 (1959), 290–297.

[12] Emilio Ferrara, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. The rise of social bots. *Commun. ACM* 59, 7 (2016), 96–104.

[13] Terrill L Frantz and Kathleen M Carley. 2017. Reporting a network's most-central actor with a confidence level. *Computational and Mathematical Organization Theory* 23, 2 (2017), 301–312.

[14] Linton C Freeman. 1977. A set of measures of centrality based on betweenness. *Sociometry* (1977), 35–41.

[15] Linton C Freeman, A Kimball Romney, and Sue C Freeman. 1987. Cognitive structure and informant accuracy. *American anthropologist* 89, 2 (1987), 310–325.

[16] Joseph Galaskiewicz. 1991. Estimating point centrality using different network sampling techniques. *Social Networks* 13, 4 (1991), 347–386.

[17] Gueorgi Kossinets. 2006. Effects of missing data in social networks. *Social networks* 28, 3 (2006), 247–268.

[18] Roy HA Lindelauf, Herbert JM Hamers, and BGM Husslage. 2013. Cooperative game theoretic centrality analysis of terrorist networks: The cases of jemaah islamiyah and al qaeda. *European Journal of Operational Research* 229, 1 (2013), 230–238.

[19] Shogo Murai and Yuichi Yoshida. 2019. Sensitivity analysis of centralities on unweighted networks. In *The World Wide Web Conference*. 1332–1342.

[20] Mark EJ Newman. 2003. The structure and function of complex networks. *SIAM review* 45, 2 (2003), 167–256.

[21] Qikai Niu, An Zeng, Ying Fan, and Zengru Di. 2015. Robustness of centrality measures against network manipulation. *Physica A: Statistical Mechanics and its Applications* 438 (2015), 124–131.

[22] Martin Olsen and Anastasios Viglas. 2014. On the approximability of the link building problem. *Theoretical Computer Science* 518 (2014), 96–116.

[23] Senni Perumal, Prithwish Basu, and Ziyu Guan. 2013. Minimizing eccentricity in composite networks via constrained edge additions. In *Military Communications Conference, MILCOM 2013-2013 IEEE*. IEEE, 1894–1899.

[24] H Prüfer. 1918. Neuer Beweis eines Satzes über Permutationen. *Archiv der Mathematik und Physik. 3. Reihe* 27 (01 1918).

[25] John Scott. 2017. *Social network analysis*. Sage.

[26] Santiago Segarra and Alejandro Ribeiro. 2015. Stability and continuity of centrality measures in weighted graphs. *IEEE Transactions on Signal Processing* 64, 3 (2015), 543–555.

[27] Marvin E Shaw. 1954. Group structure and the behavior of individuals in small groups. *The Journal of psychology* 38, 1 (1954), 139–149.

[28] Diana Stork and William D Richards. 1992. Nonrespondents in communication network studies: Problems and possibilities. *Group and Organization Management* 17, 2 (1992), 193–209.

[29] Maksim Tsvetovat and Kathleen M Carley. 2005. Generation of realistic social network datasets for testing of analysis and simulation tools. *Available at SSRN 2729296* (2005).

[30] Marcin Waniek, Tomasz P Michalak, Talal Rahwan, and Michael Wooldridge. 2017. On the construction of covert networks. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1341–1349.

[31] Marcin Waniek, Tomasz P Michalak, Michael J Wooldridge, and Talal Rahwan. 2018. Hiding individuals and communities in a social network. *Nature Human Behaviour* 2, 2 (2018), 139.

[32] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of small-world networks. *nature* 393, 6684 (1998), 440–442.