Hiding Individuals and Communities in a Social Network

Marcin Waniek
University of Warsaw

`vua@mimuw.edu.pl`

Tomasz Michalak
University of Oxford

& University of Warsaw

`tomasz.michalak@cs.ox.ac.uk`

Talal Rahwan
Masdar Institute

of Science and Technology

`trahwan@gmail.com`

Michael Wooldridge
University of Oxford

`michael.wooldridge@cs.ox.ac.uk`

Abstract

The Internet and social media have fueled enormous interest in social network analysis. New tools continue to be developed and used to analyse our personal connections, with particular emphasis on detecting communities or identifying key individuals in a social network. This raises privacy concerns that are likely to exacerbate in the future. With this in mind, we ask the question: *Can individuals or groups actively manage their connections to evade social network analysis tools?* By addressing this question, the general public may better protect their privacy, oppressed activist groups may better conceal their existence, and security agencies may better understand how terrorists escape detection. We first study how an individual can evade "network centrality" analysis without compromising his or her influence within the network. We prove that an optimal solution to this problem is hard to compute. Despite this hardness, we demonstrate that even a simple heuristic, whereby attention is restricted to the individual's immediate neighbourhood, can be surprisingly effective in practice. For instance, it could disguise Mohamed Atta's leading position within the WTC terrorist network, and that is by rewiring a strikingly-small number of connections. Next, we study how a community can increase the likelihood of being overlooked by community-detection algorithms. We propose a measure of concealment, expressing how well a community is hidden, and use it to demonstrate the effectiveness of a simple heuristic, whereby members of the community either "unfriend" certain other members, or "befriend" some non-members, in a coordinated effort to camouflage their community.

1   Introduction

The on-going process of datafication continues to turn many aspects of our lives into computerised data [23]. This data is being collected and analysed for various diverse applications by public and private institutions alike. One particular type of data that has received significant attention over the past decade concerns our social connections. To this end, a number of tools have been advocated for social network analysis, with particular emphasis on the detection of communities or the identification of key individuals within a network. For all their benefits, the widespread use of such tools raises legitimate privacy concerns. For instance, Mislove et al. [24] demonstrated how, by analysing Facebook's social network structure, as well as the attributes of some users, it is possible to infer otherwise-private information about other Facebook users.

To tackle such privacy issues, various countermeasures have been proposed, ranging from strict legal controls [1], through algorithmic solutions [15], to market-like mechanisms that allow participants to monetize their personal information [20]. However, to date only few such countermeasures have been implemented, leaving the privacy issue largely unresolved, e.g., as is evident from the very recent release of Facebook's "*Global Government Requests Report*" [2], which revealed a global increase in government requests to secretly access user data. Furthermore, it is unlikely that effective legal mechanisms will be introduced in countries with authoritarian regimes, where so-

cial networking sites and other internet content is policed, and anti-governmental blogs and activities are censored [17, 18].

Against this background, we ask the question: can individuals or communities proactively manage their social connections so that their privacy is less exposed to the workings of network analysis tools? To put it differently, can we disguise our standing in the network to escape detection? This matters because, on one hand, it assists the general public in protecting their privacy against intrusion from government and corporate interests; on the other hand, it assists counterterrorism units and law-enforcement agencies in understanding how criminals and terrorists could escape detection, especially given their increasing reliance of social-media survival strategies [27, 14]. To date, however, this fundamental question has received little attention in the literature, as most research efforts have focused on developing ever more sophisticated network analysis tools, rather than considering how to evade them.

To address the above question from an individual's viewpoint, we focus on three main centrality measures, namely *degree*, *closeness*, and *betweenness*, and study how one can avoid being highlighted by those measures without compromising his or her influence. Since, from a graph-theoretic perspective, this is fundamentally an optimization problem, we analyse its computational complexity to illuminate the theoretical limits of such capability as disguising oneself. Although we show that an optimal solution is often hard to compute, we demonstrate the effectiveness of a surprisingly simple heuristic, whereby the rewiring of social connections is restricted to the individual's immediate network neighbourhood. Specifically, it involves two actions that are already available on popular social-media platforms: (i) "unfriending" certain neighbours; (ii) introducing certain neighbours to each other.

From a group's viewpoint, we study how a community can conceal itself to increase the likelihood of being overlooked by community-detection algorithms. To this end, we propose a measure of concealment, designed to quantify the degree to which a group of individuals is hidden. Using this measure, we demonstrate the effectiveness of yet another simple heuristic, whereby members of the community either "unfriend" certain other members, or "befriend" some non-members to blend into the surrounding web of social connections.

2   The Model

This section presents the basic concepts and objectives; all formal definitions can be found in the Supporting In-

formation.

Centrality Measures: A measure of centrality reflects the importance of any given node in the network. Arguably, the standard centrality measures are: *degree*, *closeness* and *betweenness* [11]. In particular, for any given node v, the degree centrality focuses on the number of neighbours that v has (the more neighbours the better). In contrast, the closeness centrality quantifies the importance of v based on its average distance to other nodes (the closer the better). Finally, the betweenness centrality focuses on the number of shortest paths on which v lies (the more paths the better).

Models of Influence: The best established mathematical models of influence are the *Independent Cascade* model [12] and the *Linear Threshold* model [16]. Basically, both models start with some "active" subset of nodes, called the *seed set*.[1] Then, as time passes (in discrete rounds), new nodes become activated due to the influence from other previously-activated nodes. The two models differ in the way influence propagates through the network. Specifically, in the Independent Cascade model, an active node activates each of its neighbours with some pre-defined probability. In contrast, with the Linear Threshold model, each node has some random pre-defined threshold, and gets activated when the number of its active neighbors exceeds that threshold. Under either model, the influence of a node, v, on another, w, is measured as the probability that w gets activated when the seed set is $\{v\}$.

First Objective: Given a network and a *source node*, $v^\dagger$, our objective is to conceal the importance of $v^\dagger$ by decreasing its centrality (according to each of the aforementioned measures of centrality) without compromising its influence over the network (according to the aforementioned models of influence). We do so by rewiring the links of the network, without exceeding a certain *budget*—the maximum number of links allowed to be modified (i.e., added or removed). To simplify our analysis, we divide the process of disguising $v^\dagger$ into two consecutive phases. In the first phase, part of the budget is spent on minimizing the three centrality measures, during which the influence of $v^\dagger$ is likely to decrease—we call this the *centrality minimization* problem. The second phase involves spending the remaining budget to recover as much as possible of the influence of $v^\dagger$ while avoiding the addition of any links that were removed during the centrality

---

[1]An *active* node can be thought of as an *infected* person who influences, but not necessarily infects, his or her neighbours. Analogously, an *inactive* node can be a *healthy* person who is influenced by any infected neighbours he or she may have; stronger influence corresponds to stronger chances of infection.

2

minimization phase. Here, we consider two variants of this latter problem: (i) the *individual influence recovery* problem, where the goal is to recover the influence of $v^\dagger$ over every single node, and (ii) the *global influence recovery* problem, where the goal is to recover the *sum of influences* of $v^\dagger$ over all nodes.

Second Objective: Given a community, i.e., subset of nodes, $C^\dagger$, our goal is the conceal the identity of $C^\dagger$ by hiding its existence within the network. Recall that a community structure is a partition of the set of nodes into disjoint and exhaustive subsets, or "communities". As such, $C^\dagger$ is exposed if a community-detection algorithm is able to return a community structure, CS, such that $C^\dagger \in CS$. We hide $C^\dagger$ by rewiring the links of the network, again according to some budget, i.e., maximum number of permitted modifications.

3 Disguising Individuals

3.1 Hardness Results

Our main theoretical results are summarized in Table 1 (for more details, see theorems 1 through 4 in the *Supporting Information*). As shown in the table, all the problems under consideration turn out to be NP-complete, with the exception of minimizing degree centrality. To put it differently, finding an optimal way to disguise one's importance in a social network is extremely difficult (from a computational point of view), not to mention the fact that it requires knowing the entire network structure, and may also require adding or removing links that are far from the source node.

Disguising centrality (Degree)	P
Disguising centrality (Closeness)	NPC
Disguising centrality (Betweenness)	NPC
Individual influence recovery (LT)	NPC
Individual influence recovery (IC)	NPC
Global influence recovery (LT)	NPC
Global influence recovery (IC)	NPC

Table 1: Summary of our computational-hardness results.

3.2 A Scalable Heuristic

Typically, one has very limited knowledge of the social ties beyond his or her immediate friends, or maybe friends of friends. However, even if one was able to somehow acquire information about the entire network structure, our theoretical results from the previous subsection suggest

that it is extremely unlikely for such an individual to have the necessary computational power to *optimally* disguise himself or herself. Against this background, we investigate the possibility of disguising one's centrality adequately (albeit not optimally) while restricting one's attention to only his or her immediate neighbourhood, and without requiring massive computational power nor expertise in sophisticated optimization techniques. With this in mind, we propose a heuristic whose instructions are simple enough for an average user of social-networking services to understand and use, regardless of their technical background. Our heuristic, called ROAM—Remove One, Add Many—is detailed in the box below, and an illustration of how it works is presented in Figure 1.

The ROAM heuristic given a budget b:
- **Step 1:** Remove the link between the source node, $v^\dagger$, and its neighbour of choice, v_0;

- **Step 2:** Connect v_0 to $b - 1$ nodes of choice, who are neighbours of $v^\dagger$ but not of v_0 (if there are fewer than $b - 1$ such neighbours, connect v_0 to all of them).

Let us now comment on this heuristic, starting with **Step 1**. As far as the centrality of $v^\dagger$ is concerned, this step can only be beneficial. More specifically, cutting off $v^\dagger$ from one of its neighbours is the only way to reduce the degree of $v^\dagger$. Likewise, **Step 1** can only decrease the closeness of $v^\dagger$ (this happens when all shortest paths between $v^\dagger$ and some other node run through the removed link), and can only decrease the betweenness of $v^\dagger$ (this happens when some of the shortest paths going through $v^\dagger$ contain the removed link). However, as far as the influence of $v^\dagger$ is concerned, **Step 1** may be detrimental, as it deprives $v^\dagger$ from its direct influence over v_0.

Moving on to **Step 2**, this step is primarily designed to compensate for any influence that $v^\dagger$ may have lost during the previous step. Specifically, it creates new, indirect connections between $v^\dagger$ and v_0 to compensate for the direct one that was removed earlier. As far as the centrality of $v^\dagger$ is concerned, while **Step 2** does not affect the degree of $v^\dagger$, it increases the degrees of some of its neighbours, which in turn contributes towards concealing the relative importance of $v^\dagger$ within the network. Furthermore, the addition of a link, (v_0, v_i)—where v_i is some neighbour of $v^\dagger$—cannot increase the closeness centrality of $v^\dagger$ beyond its original state, i.e., its state before running the ROAM heuristic altogether. This is because any path containing (v_0, v_i) and $(v_i, v^\dagger)$ is certainly longer than an original path in which (v_0, v_i) and $(v_i, v^\dagger)$ were replaced with $(v_0, v^\dagger)$. Likewise, the addition of this link cannot
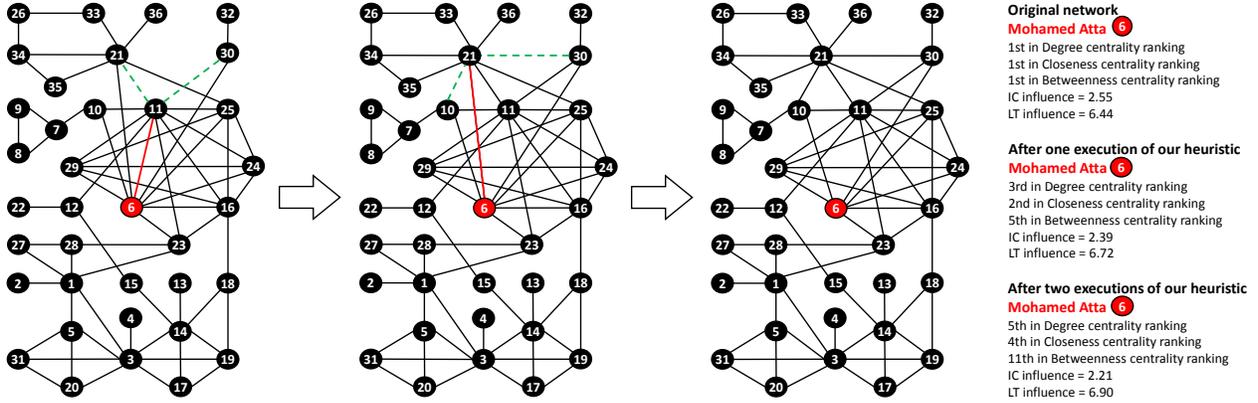
Figure 1: Executing the ROAM heuristic twice on the 9/11 terrorist network to hide Mohamed Atta—one of the ringleaders of the attack [19]. The red link is the one to be to removed by the algorithm, and the dashed links are the ones to be added.

increase the betweenness centrality of $v^\dagger$ beyond its original state, because replacing a direct connection between $v^\dagger$ and v_0 with an indirect one cannot increase the percentage of shortest paths going through $v^\dagger$.

Finally, let us comment on the how to choose v_0, and how to choose the neighbours of $v^\dagger$ to connect to v_0. Based on the simulation study reported in the Supporting Information, we choose v_0 to be the neighbour of $v^\dagger$ with the most connections, and we connect v_0 to the $b-1$ neighbours of $v^\dagger$ with the least connections. With such choices, it is relatively straightforward to execute the ROAM heuristic on existing social-networking services. On Facebook, for example, one can typically view the number of friends that each of his friends has (even if some of them make this information private, one can still choose among those that do not). Once the nodes are chosen, **Step 1** simply requires $v^\dagger$ to "unfriend" v_0, whereas **Step 2** requires $v^\dagger$ to "suggest" the friendship of v_0 to the other chosen nodes. Note that, on Facebook, $v^\dagger$ can only introduce two individuals to each other if they were both $v^\dagger$'s friends. As such, **Step 2** must be executed before **Step 1**, that is, $v^\dagger$ must end the friendship with v_0 *after* introducing v_0 to the other nodes.

4 Disguising Communities

4.1 A Measure of Concealment

We propose a measure of how well a community, $C^\dagger$, is hidden in a community structure, CS. Note that $C^\dagger$ is not necessarily a member of CS. To put it differently, when describing $C^\dagger$ as a "community", we mean to use this term in its broader sense, where $C^\dagger$ is essentially just

a subset of nodes. As such, when measuring how well $C^\dagger$ is hidden in CS, it may well be the case that the members of $C^\dagger$ are spread out across multiple communities in CS.

To this end, we start by proposing two measures, denoted by μ' and μ'', which capture different aspects of concealment. In particular, μ' is defined for every community $C^\dagger \subseteq V$ and every community structure CS as follows:

$$\mu'(C^\dagger, CS) = \frac{|\{C_i \in CS : C_i \cap C^\dagger \neq \emptyset\}| - 1}{\max(|CS| - 1, 1) \max_{C_i \in CS}(|C_i \cap C^\dagger|)}.$$

Basically, this measure focuses on how well the members of $C^\dagger$ are spread out across the communities in CS. In more detail, we have $\mu'(C^\dagger, CS) \in [0, 1]$, and the greater $\mu'(C^\dagger, CS)$, the greater the concealment of $C^\dagger$ in CS. Note that the numerator grows linearly with the number of communities that $C^\dagger$ is distributed over. Subtracting 1 from both the numerator and the $|CS|$ term of the denominator is meant to handle the worst case, where all members of $C^\dagger$ appear in a single (possibly larger) community in CS; in this case, we have: $\mu'(C^\dagger, CS) = 0$. In contrast, the term $\max_{C \in CS}(|C \cap C^\dagger|)$ is meant to promote community structures in which the members of $C^\dagger$ are more evenly distributed across the communities in CS. As such, the maximum concealment is achieved when the members of $C^\dagger$ are uniformly distributed, with each member appearing in a separate community; in this case: $\mu'(C^\dagger, CS) = 1$.

Moving on to the second measure, μ'', it is defined as:

$$\mu''(C^\dagger, CS) = \sum_{C_i \in CS} \frac{|C_i \setminus C^\dagger|}{\max(n - |C^\dagger|, 1)}.$$

Intuitively, μ'' focuses on how well $C^\dagger$ is "hidden in

4

(a) $\mu(C^\dagger, CS) = 0$ **(b)** $\mu(C^\dagger, CS) = \frac{3}{8}$ **(c)** $\mu(C^\dagger, CS) = 1$
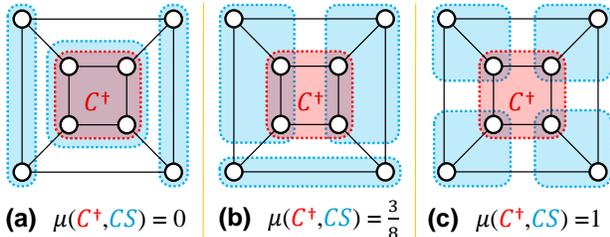
Figure 2: How the concealment of $C^\dagger$ differs from one community structure to another according to μ where $\alpha = 0.5$.

the crowd"; it grows linearly with the number of non-members of $C^\dagger$ that appear with members of $C^\dagger$ in the same community in CS. Note that $\mu''(C^\dagger, CS) \in [0, 1]$, and the greater the value, the greater the concealment of $C^\dagger$ in CS.

Having defined μ' and μ'', we now use the two as building blocks to construct a single measure whereby the trade-off between μ' and μ'' is controlled by a parameter, $\alpha \in [0, 1]$. More formally, *our proposed measure of concealment of a community $C^\dagger$ in a community structure CS is*:

$$\mu(C^\dagger, CS) = \alpha\mu'(C^\dagger, CS) + (1 - \alpha)\mu''(C^\dagger, CS).$$

Figure 2 presents a sample network with three different community structures, and highlights the community that we wish to conceal, namely $C^\dagger$. For every such community structure, we measure the concealment of $C^\dagger$ using our measure μ with $\alpha = 0.5$. In particular, Figure 2(a) presents one extreme where $\mu(C^\dagger, CS) = 0$, reflecting the fact that $C^\dagger$ is completely exposed as a community. Figure 2(b) presents the other extreme where $\mu(C^\dagger, CS) = 1$, reflecting the fact that $C^\dagger$ is completely hidden, since every member appears in a separate community along some non-member of $C^\dagger$. Finally, a case between the two extremes is presented in Figure 2(c), where $\mu(C^\dagger, CS) = \frac{3}{8}$.

4.2 A Scalable Heuristic

We set to develop a simple heuristic that can be applied by any group of people regardless of their technical background or their knowledge of the network topology. After all, it is of little use to have an exact algorithm that can only be understood or applied by optimization experts armed with enormous processing power. Likewise, exact algorithms that require knowing the entire network topology may prove useless, since such knowledge is rarely available.

Our heuristic, called DICE—Disconnect Internally, Connect Externally—is detailed in the box below.

The DICE heuristic given a budget b:
- **Step 1:** Disconnect $d \leq b$ links from within $C^\dagger$;

- **Step 2:** Connect $b - d$ nodes from within $C^\dagger$ to $b - d$ nodes from outside of $C^\dagger$.

This heuristic is inspired by *modularity* [26]—a widely used index for measuring the quality of any given community structure. Specifically, it promotes structures that have dense connections *within* communities and sparse connections *between* them. As such, community-detection algorithms are typically designed to search for a structure that maximizes modularity. With this in mind, **Step 1** of our heuristic decreases the density of the connections within $C^\dagger$, whereas **Step 2** increases the connections between $C^\dagger$ and other communities. In doing so, a community-detection algorithm is more likely to overlook $C^\dagger$, i.e., it would fail to recognize $C^\dagger$ as a community, and instead assign its members to multiple communities.

Finally, let us comment on how DICE can be applied in practice. On Facebook, for example, **Step 1** requires some members to "unfriend" other members, which is rather straightforward. As for **Step 2**, members must send a friendship request to non-members; these could be classmates, coworkers, neighbours living next door, or even random people (it is possible to try multiple random friendship requests, hoping that some of them would be successful).

5 Experiments

5.1 Data sets

We experiment with two types of real-life networks:

(a). *Covert organizations*: we consider three terrorist network, responsible for the WTC 9/11 attacks [19]; the 2002 Bali attack [13]; and the 2004 Madrid train bombings [13];

(b). *Social networks*: we study anonymized fragments of three social networks, namely Facebook, Twitter and Google+. These fragments are taken from SNAP—the Stanford Network Analysis Platform [21].

We also study randomly-generated networks, namely:

(a). *Scale-free* networks using the Barabasi-Albert model [4]. We write $ScaleFree(x, y)$ where x is the number of nodes; y is the number of links added with each node;
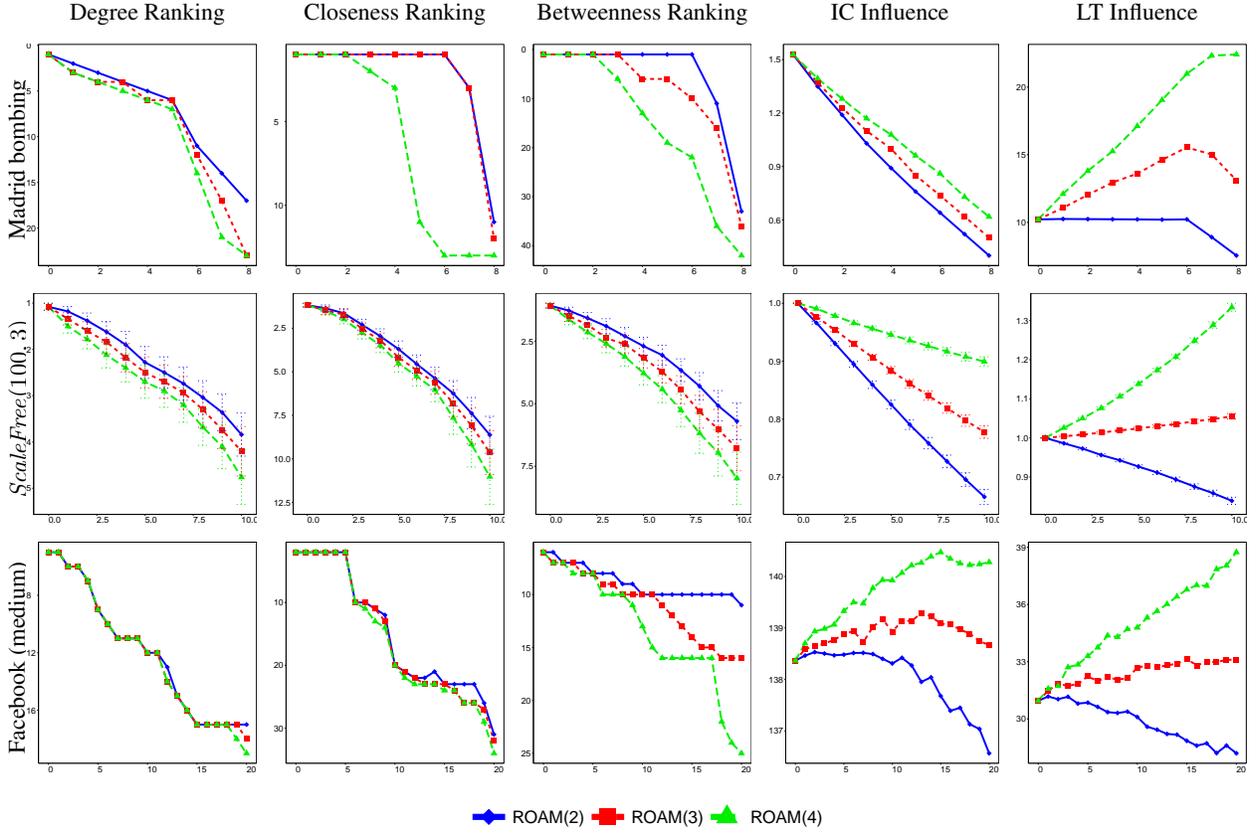
Figure 3: Executing ROAM multiple, consecutive times (the x-axis represents the number of executions). Given the Madrid-attack network, 50 scale-free networks, and a medium-sized fragment of Facebook's network (333 nodes, 5038 edges), the subfigures show the source node's ranking (according to different centrality measures), and the relative change in its influence value (according to different influence models). Results are for ROAM(b) : $b = 2, 3, 4$, where b is the budget in each execution.

(b). *Small-world* networks using the Watts-Strogatz model [33]. We write $SmallWorld(x, y, z)$ where x is the number of nodes; y is the average degree; z is the rewiring probability;

(c). *Random graphs* generated using the Erdos-Renyi model [9]. We write $RandomGraph(x, y)$ where x is the number of nodes; y is the expected average degree.

For each type of randomly-generated networks, we report the average result taken over 50 such networks, with the error bars representing the 95% confidence intervals.

5.2 Experimenting with ROAM

Each of our experiments consists of a network, a budget, a source node, and an influence model. More specifically, we experiment with a budget of 2, 3, and 4. The source node is assumed to be the one with the lowest sum of centrality rankings (ties are broken uniformly at random). Whenever the *Independent Cascade* model is used, an activation probability of 0.15 is assumed on each link. On the other hand, whenever the *Linear Threshold* model is used, a uniform distribution of thresholds is assumed (see the Supporting Information for more details). For both models, the influence values are approximated using the Monte-Carlo method. In each of these experiment, the ROAM heuristic is executed multiple, consecutive times.

Figure 3 shows the results of some of our experiments (the remaining results are provided in the Supporting In-
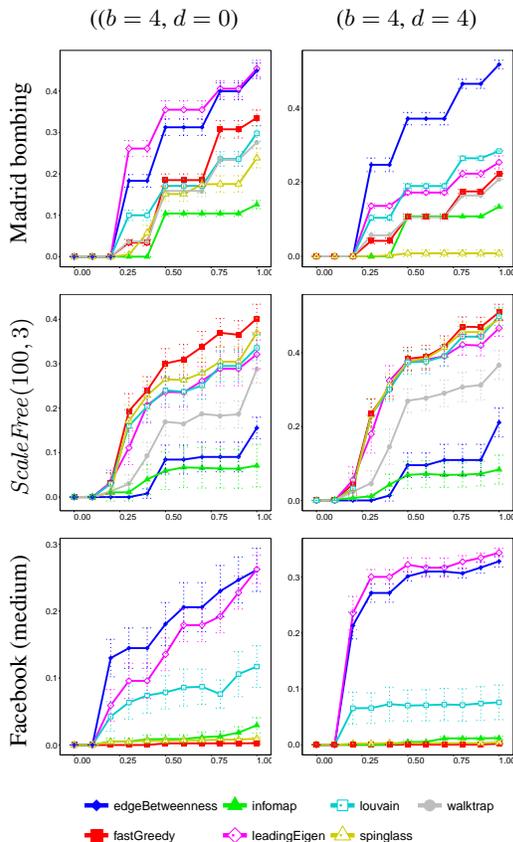
6

Figure 4: Executing DICE multiple, consecutive rounds (the x-axis represents the percentage of completed rounds), given the Madrid-attack network, scale-free networks, and a fragment of Facebook's network (consisting of 333 nodes, 5038 edges).

Figure 5: Avg. concealment-measure value in each experiment.

5.3 Experimenting with DICE

For each network, we experiment with seven community-detection algorithms implemented in the *igraph* package of the R language (version 1.0.1), namely: Eigenvector [25], Betweenness [26], Walktrap [29], Louvain [6], Greedy [7], Infomap [31] and Spinglass [30]. As such, every experiment consists of a community-detection algorithm and a network. The experiment starts by running the algorithm to obtain a community structure, CS. After that, the community to be hidden, i.e., $C^\dagger$, is chosen to be the element in CS whose size is the median of the sizes of all communities in CS (ties are broken uniformly at random). Although $C^\dagger$ does not necessary have to be an element of CS, we choose it this way in order to study the worst case scenario in which $C^\dagger$ is initially exposed completely. The experiment then proceeds in rounds, each involving the execution of DICE followed by the execution of the community-detection algorithm, to measure how well $C^\dagger$ is hidden in the new outcome of the algorithm (this measurement is done using μ with $\alpha = 0.5$). We set the number of rounds to be $\lceil |C^\dagger|/b \rceil$. In each round, we disconnect d links from within $C^\dagger$ (chosen uniformly at random), and then connect $b - d$ members of $C^\dagger$ to $b - d$ non-members of $C^\dagger$ (again chosen uniformly at random). Due to this randomness in our implementation, DICE may yield different results in different executions. Therefore, we repeat each experiment multiple times, and report the 95% confidence interval.

Figure 4 shows the results of some of our experiments (for the remaining results see the Supporting Information). As can be seen, DICE is able to hide the community, $C^\dagger$ with varying levels of success, depending on the community-detection algorithm being used. Importantly, the performance does not appear to be overly-sensitive to

formation). The centrality plots depict the *ranking* of the source node, whereas the influence plots depict its *relative* influence value (compared to the *original* influence value before executing the heuristic altogether). As can be seen, the heuristic is effective in decreasing the source node's ranking, and this effectiveness increases with the budget spent on rewiring the network. As for influence, the performance of the heuristic varies depending on the network, the influence model, and the budget. Overall, the greater the budget, the greater the influence, e.g., a budget of 4 manages to maintain (or even increase) the influence in 4 out of 6 cases.
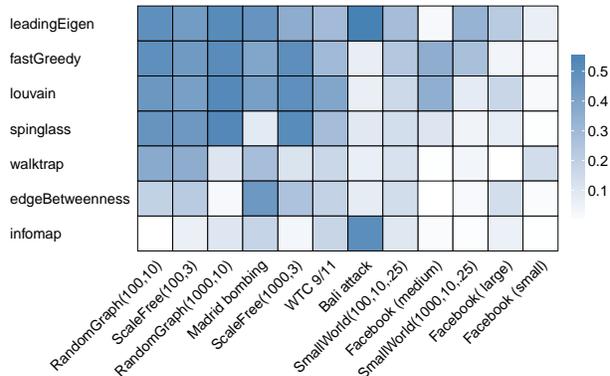
the parameter d. This is important because it provides the members of $C^\dagger$ with the ability to control this parameter as needed (i.e., control the trade-off between the number of internal edges being removed, and the number of external edges being added). For example, the members of $C^\dagger$ might be interested in hiding their community as much as possible, while removing as few internal links as possible (after all, the added external links are fake, serving no purpose other than disguising the community, whereas the removed internal links are real; they existed in the community for a reason). In such a case, since the addition of an external link is not entirely under the control of $C^\dagger$ (as it requires the consent of a non-member), the number of newly-added external links may be insufficient for providing a satisfactory level of concealment, in which case the members can compensate for this by sacrificing more internal links, i.e., by increasing the parameter d.

Figure 5 illustrates the average value of our concealment measure, μ, in each experiment where $b = 4$ and $d = 2$. In particular, each row represents a community-detection algorithm, each row represents a network, and the intensity of the colour in each cell represents the average value of μ, taken over 50 simulations, either by generating a new random network in each simulation, or by re-running the simulation over and over on the same real-life network (recall that our implementation of DICE is non-deterministic, and may yield different results on the same network). As can be seen, the Infomap [31] algorithm seems to be the most difficult to fool.

6   Discussion

Our goal was to understand the practical limits of disguising individuals and communities, to increase the likelihood of them being overlooked by social network analysis tools. Our main result is that, despite the hardness of finding an optimal solution, disguise is surprisingly easy in practice using simple heuristics that are readily-implementable even by lay people. Viewed from a different perspective, our work can be seen as an extension of the sensitivity analyses of centrality measures [8] and community detection algorithms [28]; while such analyses typically consider the effects of small network alterations, we consider changes that are much wider in scope, and strategic in nature.

On one hand, our findings contribute towards charting the limits of protecting privacy in social networks. On the other hand, they expose implications for using generic social network analysis tools in security applications; the fact that such tools can be easily misled underlines the need for developing specialized tools that account for the

nature of links and nodes in the network and not just the topology *per se*.

Despite these findings, our understanding of how to evade social network analysis tools is still limited, with many research questions yet to be answered. For instance, we still do not know how a relationship can be hidden from the eyes of link-prediction algorithms [22], or how an individual can evade detection by Eigenvector centrality—the backbone of Google's search engine.

References

[1] European data protection supervisor, meeting the challenges of big data, opinion 7/2015.

[2] https://govtrequests.facebook.com/.

[3] J. M. Anthonisse. The rush in a graph. *Amsterdam: University of Amsterdam Mathematical Centre*, 1971.

[4] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

[5] M. A. Beauchamp. An improved index of centrality. *Behavioral Science*, 10(2):161–163, 1965.

[6] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, 2008.

[7] A. Clauset, M. E. Newman, and C. Moore. Finding community structure in very large networks. *Physical review E*, 70(6):066111, 2004.

[8] C. D. Correa, T. Crnovrsanin, and K.-L. Ma. Visual reasoning about social networks using centrality sensitivity. *Visualization and Computer Graphics, IEEE Transactions on*, 18(1):106–120, 2012.

[9] P. Erdős and A. Rényi. On random graphs i. *Publ. Math. Debrecen*, 6:290–297, 1959.

[10] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, pages 35–41, 1977.

[11] L. C. Freeman. Centrality in social networks conceptual clarification. *Social networks*, 1(3):215–239, 1979.

[12] J. Goldenberg, B. Libai, and E. Muller. Using complex systems analysis to advance marketing theory development: Modeling heterogeneity effects on new product growth through stochastic cellular automata. *Academy of Marketing Science Review*, 9(3):1–18, 2001.

[13] B. Hayes. Connecting the dots can the tools of graph theory and social-network studies unravel the next big plot? *American Scientist*, 94(5):400–404, 2006.

[14] N. F. Johnson, M. Zheng, Y. Vorobyeva, A. Gabriel, H. Qi, N. Velasquez, P. Manrique, D. Johnson, E. Restrepo, C. Song, and S. Wuchty. New online ecology of adversarial aggregates: Isis and beyond. *Science*, 352(6292):1459–1463, 2016.

[15] M. Kearns, A. Roth, Z. S. Wu, and G. Yaroslavtsev. Private algorithms for the protected in social network search. *Proceedings of the National Academy of Sciences*, page 201510612, 2016.

[16] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM, 2003.

[17] G. King, J. Pan, and M. E. Roberts. How censorship in china allows government criticism but silences collective expression. *American Political Science Review*, 107(02):326–343, 2013.

[18] G. King, J. Pan, and M. E. Roberts. Reverse-engineering censorship in china: Randomized experimentation and participant observation. *Science*, 345(6199):1251722, 2014.

[19] V. Krebs. Mapping networks of terrorist cells. *Connections*, 24:43–52, 2002.

[20] J. I. Lane, V. Stodden, S. Bender, and H. Nissenbaum, editors. *Privacy, big data, and the public good: frameworks for engagement*. 2014.

[21] J. Leskovec and J. J. Mcauley. Learning to discover social circles in ego networks. In *Advances in neural information processing systems*, pages 539–547, 2012.

[22] L. Lü and T. Zhou. Link prediction in complex networks: A survey. *Physica A*, 390(6):11501170, 2011.

[23] V. Mayer-Schnberger. *Big Data: A Revolution That Will Transform How We Live, Work and Think. Viktor Mayer-Schnberger and Kenneth Cukier*. John Murray Publishers, UK, 2013.

[24] A. Mislove, B. Viswanath, K. P. Gummadi, and P. Druschel. You are who you know: Inferring user profiles in online social networks. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, WSDM '10, pages 251–260, New York, NY, USA, 2010. ACM.

[25] M. E. Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical review E*, 74(3):036104, 2006.

[26] M. E. Newman and M. Girvan. Finding and evaluating community structure in networks. *Physical review E*, 69(2):026113, 2004.

[27] A. Nordrum. Pro-ISIS Online Groups Use Social Media Survival Strategies to Evade Authorities, 2016.

[28] G. K. Orman and V. Labatut. A comparison of community detection algorithms on artificial networks. In *Discovery science*, pages 242–256. Springer, 2009.

[29] P. Pons and M. Latapy. Computing communities in large networks using random walks. In *Computer and Information Sciences-ISCIS 2005*, pages 284–293. Springer, 2005.

[30] J. Reichardt and S. Bornholdt. Statistical mechanics of community detection. *Physical Review E*, 74(1):016110, 2006.

[31] M. Rosvall, D. Axelsson, and C. T. Bergstrom. The map equation. *The European Physical Journal Special Topics*, 178(1):13–23, 2010.

[32] M. E. Shaw. Group structure and the behavior of individuals in small groups. *The Journal of Psychology*, 38(1):139–149, 1954.

[33] D. J. Watts and S. H. Strogatz. Collective dynamics of small-world networks. *nature*, 393(6684):440–442, 1998.

[34] J. Xie, S. Kelley, and B. K. Szymanski. Overlapping community detection in networks: The state-of-the-art and comparative study. *ACM Computing Surveys (csur)*, 45(4):43, 2013.

A Organization of the Appendix

In this document, we formally define the relevant centrality measures and influence models, before defining our optimization problems (Section B). After that, we present the proofs of our theoretical results (Section C), followed by a discussion of various experimental results (Section D). Finally, we study the problem of constructing a network from scratch, designed for the sole purpose of concealing the identity of the leader while ensuring that it is a highly influential node in the network (Section E).

B Definitions

Basic Notation: Let $G = (V, E) \in \mathbb{G}$ denote a network, where $V = \{v_1, \ldots, v_n\}$ is the set of n nodes and $E \subseteq V \times V$ is the set of edges. A path is a sequence of distinct nodes, $\langle v_l, \ldots, v_k \rangle$, such that every two consecutive nodes are connected by an edge. The length of a path is considered to be the number of edges in that path. For any pair of nodes, v_i, v_j in G, the set of all shortest paths between them is denoted by $sp_G(v_i, v_j)$, and the distance between them is denoted by $d_G(v_i, v_j)$, where distance is defined as the length of a shortest path between the two. In case of an *undirected* network G we do not discern between edges (v_i, v_j) and (v_j, v_i); otherwise the network is said to be *directed*. Furthermore, G is said to be *connected* (*strongly connected* for directed networks) if there exists a path between every pair of nodes in G.

We denote by $N_G^{pred}(v_i)$ the set of *predecessors* of v_i in G, that is, $N_G^{pred}(v_i) = \{v_j \in V : (v_j, v_i) \in E\}$. On the other hand, we denote by $N_G^{succ}(v_i)$ the set of *successors* of v_i in G, i.e., $N_G^{succ}(v_i) = \{v_j \in V : (v_i, v_j) \in E\}$. Finally, we denote by $N_G(v_i)$ the set of neighbours of v_i in G, i.e., $N_G(v_i) = N_G^{pred}(v_i) \cup N_G^{succ}(v_i)$. For the case of undirected graph, we will assume that $N_G(v_i) = N_G^{pred}(v_i) = N_G^{succ}(v_i)$.

To make the notation more readable, we will often denote two arbitrary nodes by v and w, instead of v_i and v_j. Moreover, we will often omit the network itself from the notation whenever it is clear from the context, e.g., by writing $d(v, w)$ instead of $d_G(v, w)$; this applies not only to the notation presented thus far, but to all notation.

We consider a *community structure*, $CS = \{C_1, \ldots, C_k\}$, to be a partition of the set of nodes into disjoint and exhaustive subsets, or communities.[2] Formally, it satisfies the following three conditions: $\forall_{C_i \in CS} C_i \subseteq V$, $\bigcup_{C_i \in CS} C_i = V$, and $\forall_{C_i, C_j \in CS} C_i \cap C_j = \emptyset$.

Centrality Measures: Formally, a centrality measure [11] is a function $c : \mathbb{G} \times V \to \mathbb{R}$. The *degree* centrality [32] is denoted by c_{degr}, the *closeness* centrality [5] is denoted by c_{clos}, and the *betweeness* centrality [3, 10] is denoted by c_{betw}. Specifically, given a node $v_i \in V$ and an undirected network, we have:

$$c_{degr}(G, v_i) = \frac{|N_G(v_i)|}{n - 1}$$

$$c_{clos}(G, v_i) = \frac{n - 1}{\sum_{v_j \in V} d_G(v_i, v_j)}$$

$$c_{betw}(G, v_i) = \frac{2}{(n-1)(n-2)} \sum_{v_j, v_k \in V \setminus \{v_i\}} \frac{|\{p \in sp_G(v_j, v_k) : v_i \in p\}|}{|sp_G(v_j, v_k)|}$$

On the other hand, given a directed network, we have:

$$c_{degr}(G, v_i) = \frac{|N_G(v_i)|}{2(n - 1)}$$

$$c_{clos}(G, v_i) = \frac{1}{n - 1} \sum_{v_j \in V} \frac{1}{d_G(v_i, v_j)}$$

---

[2]Some works have considered overlapping community structures [34]. However, as common in the literature, we restrict our attention to disjoint communities.

$$c_{betw}(G, v_i) = \frac{1}{(n-1)(n-2)} \sum_{v_j, v_k \in V \setminus \{v_i\}} \frac{|\{p \in sp_G(v_j, v_k) : v_i \in p\}|}{|sp_G(v_j, v_k)|} + \frac{|\{p \in sp_G(v_k, v_j) : v_i \in p\}|}{|sp_G(v_k, v_j)|}$$

Models of Influence: The propagation of influence through the network is often modeled as follows: when a certain node is sufficiently influenced by its neighbour(s), it becomes "active", in which case it starts to influence any "inactive" neighbour(s) it may have, and so on. Of course, to initiate this propagation process, a set of nodes needs to be activated right from the start; this set is called the *seed set*. Assuming that time moves in discrete rounds, we denote by $I(t) \subseteq V$ the set of nodes that are active at round t, implying that $I(1)$ is the seed set. The way influence propagates from the seed set to the remaining nodes depends on the influence model under consideration. Here, the two main models of influence are:

- *Independent Cascade* [12]: In this model, every pair of nodes is assigned an activation probability, $p : V \times V \to [0, 1]$. Then, in every round, $t > 1$, every node $v \in V$ that became active in round $t - 1$ activates every inactive successor, $w \in N^{succ}(v) \setminus I(t-1)$, with probability $p(v, w)$. The process ends when there are no new active nodes, i.e., when $I(t) = I(t-1)$.

- *Linear Threshold* [16]: In this model, every node $v \in V$ is assigned a *threshold value* t_v which is sampled (according to some probability distribution) from the set $\{0, \ldots, |N^{pred}(v)|\}$. Then, in every round, $t > 1$, every inactive node v becomes active, i.e., becomes a member of $I(t)$, if: $|I(t-1) \cap N^{pred}(v)| \geq t_v$. The process ends when $I(t) = I(t-1)$.

In either model, the influence of a node, v, on another, w, is denoted by $inf_G(v, w)$ and defined as *the probability that w gets activated given the seed set* $\{v\}$ (we make the common assumption that $inf_G(v, v) = 0$ for all $v \in V$). The influence of v over the entire network G is then: $inf_G(v) = \sum_{w \in V} inf_G(v, w)$.

First Objective (Disguising a Node): Roughly speaking, given a *source node*, $v^\dagger$, and a limited budget, b, specifying the maximum number of edges that are allowed to be added or removed, our goal is to first rewire the network so as to minimize the centrality of $v^\dagger$, and then to further rewire the network so as to "recover" the influence of $v^\dagger$ (in an attempt to compensate for any influence that $v^\dagger$ might have lost during the centrality-minimization phase). We consider two variants of the influence-recovery problem; the first focuses on the influence of $v^\dagger$ over every single node, whereas the second focuses on the influence of $v^\dagger$ over the network as a whole. In both cases, only the addition of edges is considered, since the removal of edges can only decrease the influence of $v^\dagger$. Next, we formally define the aforementioned problems.

Definition 1 (Disguising Centrality) *This problem is defined by a tuple, $(G, v^\dagger, b, c, \hat{R}, \hat{A})$, where $G = (V, E) \in \mathbb{G}$ is a network, $v^\dagger \in V$ is the source node (whose centrality is to be minimized), $b \in \mathbb{N}$ is a budget specifying the maximum number of edges that can be added or removed, $c : \mathbb{G} \times V \to \mathbb{R}$ is a centrality measure, $\hat{R} \subseteq E$ is a set of edges whose removal is forbidden, $\hat{A} \subseteq (V \times V) \setminus E$ is a set of edges whose addition is forbidden. The goal is then to identify two sets of edges, $R^* \subseteq (E \setminus \hat{R})$ and $A^* \subseteq (V \times V) \setminus (E \cup \hat{A})$, such that: $|A^*| + |R^*| \leq b$ and $G^* = (V, (E \cup A^*) \setminus R^*)$ is connected (strongly connected if G is directed) and G^* is in:*

$$\underset{G' \in \left\{ (V, (E \cup A) \setminus R) : R \subseteq (E \setminus \hat{R}), A \subseteq (V \times V) \setminus (E \cup \hat{A}) \right\}}{\arg \min} c(G', v^\dagger).$$

Definition 2 (Individual Influence recovery) *This problem is defined by a tuple, $(G, v^\dagger, inf, \hat{A}, f)$, where $G = (V, E) \in \mathbb{G}$ is a network, $v^\dagger \in V$ is the source node (whose influence is to be recovered), $inf : V \times V \to \mathbb{R}$ is an influence measure, $\hat{A} \subseteq (V \times V) \setminus E$ is a set of edges whose addition is forbidden, and $f : V \to \mathbb{R}$ specifies the influences to be recovered (i.e., for every $v_i \in V$ we want the influence of $v^\dagger$ over v_i to be at least $f(v_i)$). The goal is then to identify a set of edges, A^*, that is in:*

$$\underset{A \subseteq (V \times V) \setminus (E \cup \hat{A}) : \forall_{v_i \in V} inf_{(V, E \cup A)}(v^\dagger, v_i) \geq f(v_i)}{\arg \min} |A|.$$

Definition 3 (Global Influence recovery) *This problem is defined by a tuple, $(G, v^\dagger, inf, \hat{A}, \phi)$, where $G = (V, E) \in \mathbb{G}$ is a network, $v^\dagger \in V$ is the source node (whose influence is to be recovered), $inf : V \times V \rightarrow \mathbb{R}$ is an influence measure, $\hat{A} \subseteq (V \times V) \setminus E$ is a set of edges whose addition is forbidden, and $\phi \in \mathbb{R}$ is the total influence to be recovered. The goal is then to identify a set of edges, A^*, that is in:*

$$\underset{A \subseteq (V \times V) \setminus (E \cup \hat{A}) : inf_{(V, E \cup A)}(v^\dagger) \geq \phi}{\arg\min} |A|.$$

Second Objectives (Disguising a Community): Roughly speaking, given a community to be hidden, $C^\dagger$, and a limited budget, b, specifying the maximum number of edges that are allowed to be added or removed, our goal is to rewire the network so as to hide $C^\dagger$. To this end, we propose a measure of concealment, μ, defined for every community $C^\dagger \subseteq V$ and every community structure CS, as follows:[3]

$$\mu(C^\dagger, CS) = \alpha \mu'(C^\dagger, CS) + (1 - \alpha)\mu''(C^\dagger, CS),$$

where $\alpha \in [0, 1]$ and:

$$\mu'(C^\dagger, CS) = \frac{|\{C_i \in CS : C_i \cap C^\dagger \neq \emptyset\}| - 1}{\max(|CS| - 1, 1) \max_{C_i \in CS}(|C_i \cap C^\dagger|)}$$

$$\mu''(C^\dagger, CS) = \sum_{C_i \in CS} \frac{|C_i \setminus C^\dagger|}{\max(n - |C^\dagger|, 1)}.$$

Note that $\mu(C^\dagger, CS) \in [0, 1]$ for all $C^\dagger$ and CS, with greater values indicating greater levels of concealment of $C^\dagger$ in CS. Having presented our concealment measure, we are now ready to formally introduce our problem.

Definition 4 (Disguising a Community) *This problem is defined by a tuple, $(G, C^\dagger, alg, b)$, where $G = (V, E)$ is a network, $C^\dagger \subseteq V$ is the community to be hidden, alg is a community-detection algorithm, and $b \in \mathbb{N}$ is a budget specifying the maximum number of edges that can be added or removed. The goal is then to find a set of edges to be added, $A^* \subseteq (V \times V) \setminus E$, and another to be removed, $R^* \subseteq E$, such that $|A^*| + |R^*| \leq b$ and $G^* = (V, (E \cup A^*) \setminus R^*)$ is in:*

$$\underset{G' \in \{(V, (E \cup A) \setminus R) \,:\, A \subseteq (V \times V) \setminus E, \, R \subseteq E, \, |A| + |R| \leq b\}}{\arg\max} \mu(C^\dagger, alg(G')),$$

where $alg(G)$ is the community structure returned by the algorithm alg given the network G.

Note that the above optimization problem requires $C^\dagger$ to know the exact community-detection algorithm that the adversary is using. Since such knowledge is hardly available, we avoid this requirement, and instead aim to develop a general-purpose heuristic, designed for no particular community-detection algorithm.

C   Proofs

From the computational point of view, disguising the degree centrality of $v^\dagger$ is easy, since the only way to decrease this centrality is to remove edges connecting $v^\dagger$ to its neighbour(s). Next, we study the problems of disguising closeness centrality and betweenness centrality, followed by the problem of influence recovery under the Independent-Cascade model and under the Linear-Threshold model.

Theorem 1 *Disguising closeness centrality is NP-complete.*

---

[3]Note that $C^\dagger$ is not necessarily a member of CS. To put it differently, when describing $C^\dagger$ as a "community" we use this term in its broader sense, where $C^\dagger$ is essentially just a subset of nodes. As such, when measuring how well $C^\dagger$ is hidden in CS, it may well be the case that the members of $C^\dagger$ are spread out across multiple communities in CS.

Proof. The decision version of the optimization problem is the following: given a network $G = (V, E)$, a source node $v^\dagger$, two sets $\hat{R} \subseteq E$, $\hat{A} \subseteq (V \times V) \setminus E$, a budget $b \in \mathbb{N}$ and a value $x \in \mathbb{R}$, does there exist two sets $R^* \subseteq (E \setminus \hat{R})$ and $A^* \subseteq (V \times V) \setminus (E \cup \hat{A})$ such that $|A^*| + |R^*| \leq b$, and the network $(V, (E \cup A^*) \setminus R^*)$ is connected (strongly connected if G is directed) and $c_{clos}((V, (E \cup A^*) \setminus R^*), v^\dagger) \leq x$?

This problem is in NP, as given a solution, i.e., two sets A^* and R^*, we can verify whether $c_{clos}((V, (E \cup A^*) \setminus R^*), v^\dagger) \leq x$ in polynomial time; this only requires computing the closeness centrality of node $v^\dagger$ in network $(V, (E \cup A^*) \setminus R^*)$.

We will now show that the decision version is NP-hard. To this end, let us denote by $q \in \mathbb{R}$ the smallest possible closeness centrality of $v^\dagger$ in any (strongly) connected network whose set of nodes is V. One can see that $q = 2/n$ in the case of undirected networks, and $q = (\sum_{i=1}^{n-1} \frac{1}{i})/(n-1)$ in the case of directed network; this happens if and only if:

- the network is a path of which $v^\dagger$ is an end (when dealing with undirected networks); or

- the network is a directed cycle (when dealing with directed networks).

Let us denote such a network by Q; the closeness centrality of $v^\dagger$ in Q is then q. With this in mind, the proof involves a reduction from the *Hamiltonian cycle* problem (i.e., the problem of determining whether there exists a cycle that visits each node exactly once) to the decision problem of determining whether it is possible to reduce the closeness centrality of $v^\dagger$ to a value smaller than, or equal to, q.

To this end, given some arbitrary network, $G' = (V', E')$, be it directed or undirected, let us modify G' so as to obtain a new network, $G = (V, E)$, as illustrated in figures 6 and 7. Formally, we do so by choosing some arbitrary node, $w \in V'$, and then setting:

$$V = V' \cup \{v^\dagger, v_1, v_2\}, \quad E = E' \cup \{(v^\dagger, w), (v_1, v_2)\} \cup \{(v, v_1) : v \in V', v \in N_{G'}(w)\},$$

in the case of undirected networks, or setting:

$$V = V' \cup \{v^\dagger, v_1\}, \quad E = E' \cup \{(v^\dagger, w), (v_1, v^\dagger)\} \cup \{(v, v_1) : v \in V', v \in N_{G'}^{pred}(w)\},$$

in the case of directed networks.

We will now show that the Hamiltonian cycle problem in G' is equivalent to the following decision problem: Given network G and budget $b = |E'| - |V'| + N_{G'}^{pred}(w)|$, where $\hat{A} = \hat{R} = \emptyset$, determine whether it is possible to reduce the closeness centrality of $v^\dagger$ to a value $\leq q$, by removing at most b edges from G. Throughout the remainder of the proof, the edges and nodes in G that were in G' will be referred to as "*original*".

Firstly, we will show that if G' has a Hamiltonian cycle then it is possible to obtain Q by removing $|E'| - |V'| + |N_{G'}^{pred}(w)|$ edges from G. To this end, fix a Hamiltonian cycle of G', then:

- remove from G all original edges that are not in the Hamiltonian cycle; there are exactly $|E'| - |V'|$ such edges;

- in the Hamiltonian cycle, there are exactly two edges of which w is an end; remove any of those edges in the undirected network, or the one pointing to w in the directed network; let us denote the removed edge as (v', w);

- remove all edges from all predecessors of w to v_1, with the exception of (v', v_1); there are exactly $|N_{G'}^{pred}(w)| - 1$ such edges.

In so doing, we have obtained the network Q by removing a total of $|E'| - |V'| + |N_{G'}^{pred}(w)|$ edges from G (see figures 6 and 7).

Secondly, we show that if it is possible to obtain Q by removing $|E'| - |V'| + |N_{G'}^{pred}(w)|$ edges from G, then there exists a Hamiltonian cycle in G'. We will first deal with the undirected case, before moving on to the directed case.

In the undirected case, observe that nodes $v^\dagger$ and v_2 each have a degree of 1 in G, since their only neighbours are w and v_1, respectively. Now since Q is connected, and since we obtained Q by only removing (rather than adding) edges from G, the nodes $v^\dagger$ and v_2 must each have a degree of 1 in Q. Consequently, they must be the two ends of Q. This, in turn, implies that v_1 must have exactly two neighbours in Q; we know that one of them is v_2, let us call the other v'. This, as well as the fact that $v^\dagger$ is only connected to w, implies that the segment of Q between w and v' contains all
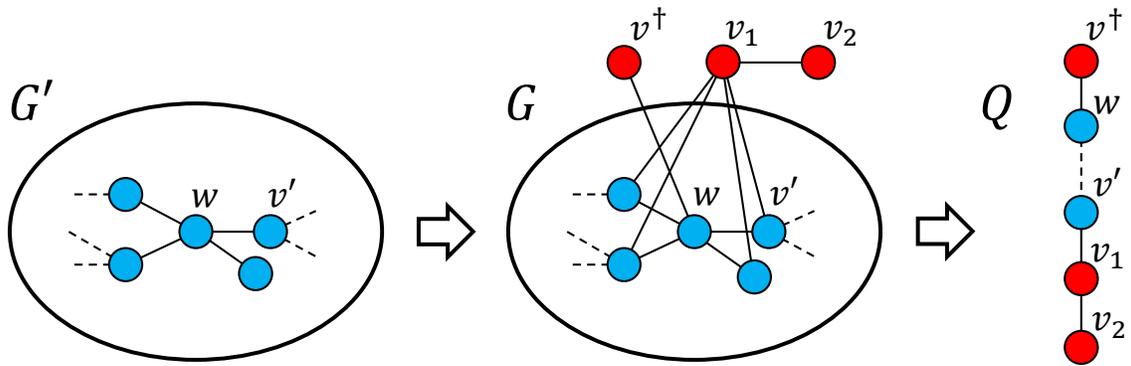
13

Figure 6: The main steps of reducing the Hamiltonian cycle problem to the problem of determining whether the closeness centrality of $v^\dagger$ can be reduced to a value $\leq q$ in an **undirected network**.
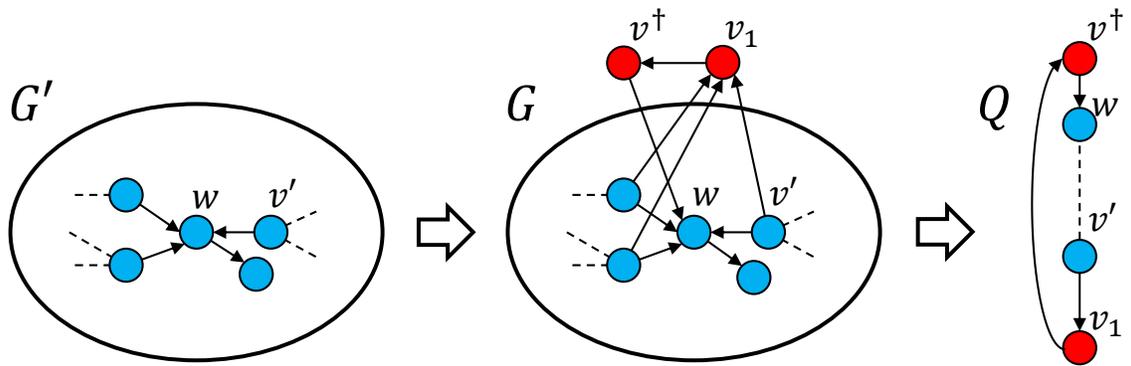

Figure 7: The main steps of reducing the Hamiltonian cycle problem to the problem of determining whether the closeness centrality of $v^\dagger$ can be reduced to a value $\leq q$ in a **directed network**.

original nodes from G' and only original edges from G' (recall that we did not add any edges between original nodes). Finally, by adding to that segment the original edge between v' and w, we obtain a Hamiltonian cycle in G'.

As for the directed case, we observe that node $v^\dagger$ has only one successor in Q, namely w, and only one predecessor in Q, namely v_1. We also know that v_1 has only one predecessor in Q; let us call that predecessor v'. These facts imply that the segment of Q between w and v' contains all original nodes from G' and only original edges from G' (again, recall that we did not add any edges between original nodes). By adding to that segment the original edge between v' and w, we obtain a Hamiltonian cycle in G'.

We have shown that a Hamiltonian cycle in G' exists if and only if it is possible to reduce the closeness centrality of $v^\dagger$ to q by removing exactly $|E'|-|V'|+|N_{G'}^{pred}(w)|$ edges from G, which concludes the proof. $\qquad\Box$

Theorem 2 *Disguising betweenness centrality is NP-complete.*

Proof. The decision version of the optimization problem is the following: given a network $G = (V, E)$, a source node $v^\dagger$, two sets $\hat{R} \subseteq E$, $\hat{A} \subseteq (V \times V) \setminus E$, a budget $b \in \mathbb{N}$ and a value $x \in \mathbb{R}$, does there exist two sets $R^* \subseteq (E \setminus \hat{R})$ and $A^* \subseteq (V \times V) \setminus (E \cup \hat{A})$ such that $|A^*|+|R^*| \le b$, and the network $(V, (E \cup A^*) \setminus R^*)$ is connected (strongly connected if G is directed) and $c_{betw}((V, (E \cup A^*) \setminus R^*), v^\dagger) \le x$?

This problem is in NP, as given a solution, i.e., two sets A^* and R^*, we can verify whether $c_{betw}((V, (E \cup A^*) \setminus R^*), v^\dagger) \le x$ in polynomial time; this only requires computing the betweenness centrality of node $v^\dagger$ in network $(V, (E \cup A^*) \setminus R^*)$.

We will now show that the decision version is NP-hard. To this end, we propose a reduction from the NP-complete *Set cover* problem. The decision version of this problem is defined by a universe $U = \{u_1, \ldots, u_l\}$ and a collection of sets $S = \{S_1, \ldots, S_m\}$ such that $\forall_j S_j \subset U$, where the goal is to determine whether there exist $k \le m$ elements of S the union of which equals U.
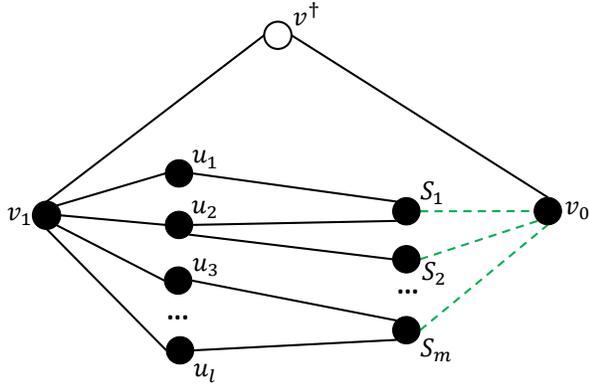

Figure 8: Undirected network used to reduce the *Set cover* problem to our problem of disguising the betweenness centrality of $v^\dagger$. To solve both problems, we consider adding (some of) the dashed edges.
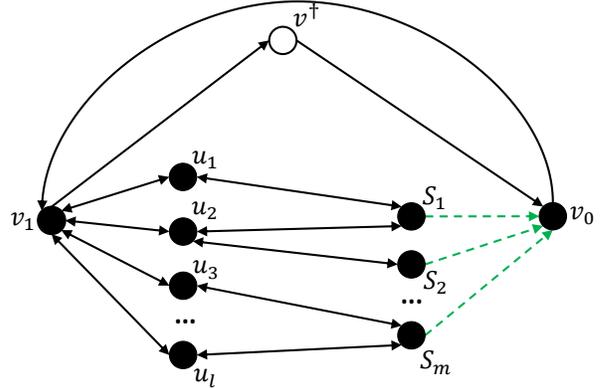
Figure 9: A directed network used to reduce the *Set cover* problem to our problem of disguising the betweenness centrality of $v^\dagger$. To solve both problems, we consider adding (some of) the dashed edges.

First, let us create a network G as shown in figures 8 and 9. More specifically, we create one node for every $S_j \in S$, one node for every $u_i \in U$, and three additional nodes, $v^\dagger$, v_0 and v_1. Next, we add (either undirected or directed) edges as follows. We add the edges $(v^\dagger, v_0)$ and $(v_1, v^\dagger)$, and for every node $u_i \in S_j$ we add the edges (S_j, u_i) and (u_i, v_1). In case of a directed network, we also add the edges (u_i, S_j) and (v_1, u_i) for every $u_i \in S_j$, as well as the edge (v_0, v_1).

Now, consider the problem of disguising the betweenness centrality of $v^\dagger$ in G given $\hat{R} = E$ and $\hat{A} = (V \times V) \setminus \{(S_1, v_0), \ldots, (S_m, v_0)\}$. Note that $v^\dagger$ "controls" (i.e., lies on) every shortest path to v_0, and does not control any shortest path between any other pair of nodes. As such, to minimize the betweenness centrality of $v^\dagger$, we need to create alternative shortest paths to v_0; this should be done by adding (some of) the edges in $\{(S_1, v_0), \ldots, (S_m, v_0)\}$,

15

since no other edge can be added, and no edge can be removed (following the definitions of $\hat{R}$ and $\hat{A}$). To be more precise, we can add at most b edges $\{(S_1, v_0), \ldots, (S_m, v_0)\}$, since we cannot exceed the budget. After this process, the betweenness centrality of $v^\dagger$ may drop to as little as $q = \frac{2}{(n-1)(n-2)}$ in the undirected case, or as little as $q = \frac{1}{(n-1)(n-2)}$ in the directed case; this happens when $v^\dagger$ no longer controls any of the shortest paths to v_0 except for the one from v_1 to v_0. Note that adding an edge (S_j, v_0) creates a new shortest path from every nodes $u_i \in S_j$ to v_0. This implies that the betweenness centrality of $v^\dagger$ can be reduced to q if and only if there exists at most b elements of S the union of which equals U.

We have just reduced the decision version of the Set Cover problem given k to the following decision problem: Given network G and budget $b = k$, where $\hat{R} = E$ and $\hat{A} = (V \times V) \setminus \{(S_1, v_0), \ldots, (S_m, v_0)\}$, determine whether it is possible to reduce the closeness centrality of $v^\dagger$ to some value $\leq q$, by removing at most b edges from G. $\qquad\square$

Theorem 3 *Both the global and the individual influence recovery problems are NP-hard under the Independent Cascade model.*

Proof. We show a reduction from the NP-complete *Set cover* problem, defined by a universe $U = \{u_1, \ldots, u_l\}$ and a collection of sets $S = \{S_1, \ldots, S_m\}$ such that $S_1 \cup \ldots \cup S_m = U$ and $\forall_j S_j \subseteq U$, and the goal is to determine whether there exist $k \leq m$ elements of S the union of which equals U.

To this end, let us create a network G as illustrated in figures 10 and 11. In more detail, we start by creating one node for every $S_j \in S$, one node for every $u_i \in U$, and one additional node $v^\dagger$. After that, for every $S_j \in S$ and every $u_i \in S_j$, we add the edge (S_j, u_i) (either directed or undirected). In the directed case we additionally add an edge $(u_i, v^\dagger)$ for every $u_i \in U$.
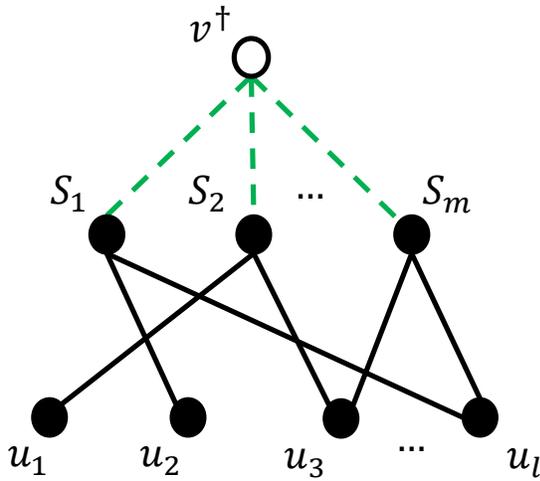

Figure 10: Undirected network used to reduce the *Set cover* problem to our influence recovery problem. To solve both problems, we consider adding (some of) the dashed edges.
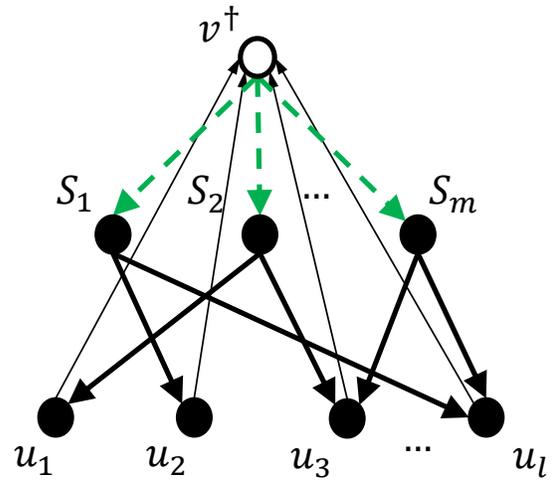

Figure 11: A directed network used to reduce the *Set cover* problem to our influence recovery problem. To solve both problems, we consider adding (some of) the dashed edges.

Consider the influence recovery problem in G under the Independent Cascade model, where:

- $\hat{A} = (V \times V) \setminus \{(v^\dagger, S_1), \ldots, (v^\dagger, S_m)\}$;

- $p : V \times V \to [0, 1]$ such that $\forall_{S_j \in S}\, p(v^\dagger, S_j) = 1$ and $\forall_{S_j \in S} \forall_{u_i \in S_j}\, p(S_j, u_i) = 1$, and $p(v, w) = 0$ for every other pair of nodes;

- for *individual* influence recovery, $\forall_{u_i \in U} f(u_i) = 1$ and $f(v) = 0$ for every other node;

- for *global* influence recovery, $\phi = k + l$.

16

The goal is then to identify the smallest subset of edges to be added to the network, $A \subseteq \{(v^\dagger, S_1), \ldots, (v^\dagger, S_m)\}$, such that either $inf_{(V, E \cup A)}(v^\dagger) \geq \phi$ in the *global* variant of the problem, or $\forall_{v_i \in V} inf_{(V, E \cup A)}(v^\dagger, v_i) \geq f(v_i)$ in the *individual* variant of the problem.

Recall that the influence of $v^\dagger$ is measured by setting the seed set as $\{v^\dagger\}$ and calculating the probability that other nodes get activated. Also recall that under the Independent Cascade model an active node, v, activates any of its predecessors, w, with probability $p(v, w)$. Importantly, with the p function defined as above, adding an edge $(v^\dagger, S_j)$ for some $S_i \in S$ makes the influence of $v^\dagger$ on every $u_i \in S_j$ equal to 1. Furthermore, the above definitions of ϕ and f imply that our goal (in both the individual and the global variants of the problem) is achieved if and only if the influence of $v^\dagger$ on *every node* $u_i \in U$ equals 1. Consequently, our goal is achieved if and only if we add to G a set of edges, $A \subseteq \{(v^\dagger, S_1), \ldots, (v^\dagger, S_m)\}$, such that:

$$\bigcup_{(v^\dagger, S_j) \in A} S_j = U.$$

Since we are interested in finding the smallest such subset, a solution to the above instance of the influence recovery problem gives us a solution to the Set Cover problem. □

Theorem 4 *Both the global and the individual influence recovery problems are NP-hard under the Linear Threshold model.*

Proof. We show a reduction from the NP-complete *Set cover* problem, defined by a universe $U = \{u_1, \ldots, u_l\}$ and a collection of sets $S = \{S_1, \ldots, S_m\}$ such that $S_1 \cup \ldots \cup S_m = U$ and $\forall_j S_j \subseteq U$, and the goal is to determine whether there exist $k \leq m$ elements of S the union of which equals U.
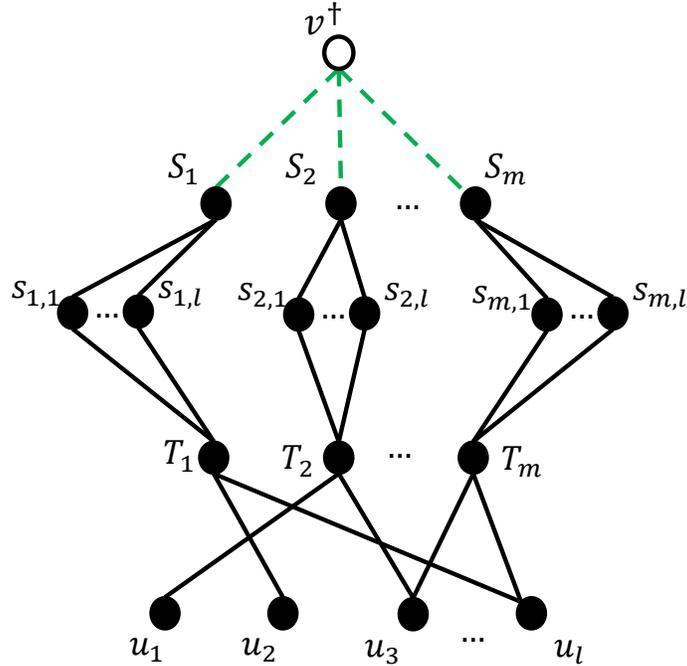

Figure 12: Undirected network used to reduce the *Set cover* problem to our influence recovery problem under the Linear Threshold model. To solve both problems, we consider adding (some of) the dashed edges.

For the directed case, we create a network G as illustrated earlier in Figure 11. As for the undirected case, we create G as illustrated in Figure 12. In more detail, for every $S_j \in S$, we create two nodes, namely S_j and T_j, as well as l

additional nodes, namely $s_{j,1}, \ldots, s_{j,l}$. We also create one node for every $u_i \in U$, and finally add the source node, $v^\dagger$. As for the edges, for every $S_j \in S$ and every $u_i \in S_j$, we add the edge (T_j, u_i). Furthermore, for every node $s_{j,i}$, we add the edges $(S_j, s_{j,i})$ and $(s_{j,i}, T_j)$.

Now consider the influence recovery problem in G under the Linear Threshold model, where:

- $\hat{A} = (V \times V) \setminus \{(v^\dagger, S_1), \ldots, (v^\dagger, S_m)\}$;

- $t_v = l$ for every node $v \in \{T_1, \ldots, T_m\}$ and $t_v = 1$ for every other node;

- for *individual* influence recovery, $\forall_{u_i \in U} f(u_i) = 1$ and $f(v) = 0$ for every other node;

- for *global* influence recovery, $\phi = k + l$ for the directed case, and $\phi = k(l + 2) + l$ for the undirected case.

The goal is then to identify the smallest subset of edges to be added to the network, $A \subseteq \{(v^\dagger, S_1), \ldots, (v^\dagger, S_m)\}$, such that either $inf_{(V,E \cup A)}(v^\dagger) \geq \phi$ in the *global* variant of the problem, or $\forall_{v_i \in V} inf_{(V,E \cup A)}(v^\dagger, v_i) \geq f(v_i)$ in the *individual* variant of the problem.

Recall that the influence of $v^\dagger$ is measured by setting the seed set as $\{v^\dagger\}$ and calculating the probability that other nodes get activated. Also recall that under the Linear Threshold model a node, v, gets activated if the number of its active predecessors exceeds t_v. Note that, with t_v defined as above, adding an edge $(v^\dagger, S_j)$ in the undirected case leads to the activation of nodes $s_{i,j}$ and T_i, which in turn leads to the activation of every $u_i \in S_j$ (see Figure 12). Likewise, in the directed case, adding $(v^\dagger, S_j)$ leads to the activation of every $u_i \in S_j$ (see Figure 11). To put it differently, when adding $(v^\dagger, S_j)$, the influence of $v^\dagger$ on every $u_i \in S_j$ equals 1. Importantly, the above definitions of ϕ and f imply that our goal (in both the individual and the global variants of the problem) is achieved if and only if the influence of $v^\dagger$ on *every node $u_i \in U$* equals 1. Those observations imply that our goal is achieved if and only if we add to G a set of edges, $A \subseteq \{(v^\dagger, S_1), \ldots, (v^\dagger, S_m)\}$, such that:

$$\bigcup_{(v^\dagger, S_j) \in A} S_j = U.$$

Since we are interested in finding the smallest such subset, a solution to the above instance of the influence recovery problem gives us a solution to the Set Cover problem. □

D   Empirical Evaluation

D.1   Configuring the ROAM Heuristic

As mentioned in the main article, the ROAM heuristic involves choosing v_0 (the neighbour of $v^\dagger$ whom the heuristic will disconnect from $v^\dagger$), and choosing the $b-1$ neighbours of $v^\dagger$ whom the heuristic will connect to v_0. We conducted a number of experiments to determine whether it is more beneficial to choose v_0 as the neighbour of $v^\dagger$ with the *least* connections or the *most* connections. Likewise, we wanted to determine whether it is more beneficial to choose the $b-1$ neighbours of $v^\dagger$ (who will be connected to v_0) as the ones with the *least* connections or the *most* connections. In particular, Figure 13 compares the different settings given 50 radomly generated scale-free networks consisting of 100 nodes each, where 3 edges are added with each step of the generation process (for more details, see [4]); we chose scale-free networks as they resemble real-life networks in many way, e.g., in terms of degree distribution. As for the source node, it is chosen to be the one with the lowest sum of centrality rankings (ties are broken uniformly at random). As for the Independent Cascade model, we set the activation probability to be $p(v, w) = 0.15$ for every pair of nodes, $v, w \in V$. As for the Linear Threshold model, for every node, $v \in V$, the threshold value, t_v, is sampled uniformly at random from the set $\{0, \ldots, |N^{pred}(v)|\}$. For both models, the influence values are approximated using the Monte-Carlo method. In the figure, we write ROAM-x-$y(b)$, where x can either be "*max*" or "*min*" (indicating that v_0 is the neighbour with the *most* connections or the *least* connections, respectively) and y can either be "*max*" or "*min*" (indicating that the $b-1$ neighbours are chosen to be the ones with the *most* connections or the *least* connections, respectively), whereas b represents the budget (which is set to 3 in this experiment). Since the results are averaged over

50 random networks, the error bars in the figure represent the 95% confidence intervals. For each network, the ROAM heuristic is executed multiple, consecutive times; the x-axis in each subfigure represents the number of executions. As can be seen, while there is no setting that dominates the others, the best overall performance seems to be achieved by ROAM-max-min(3). Based on this, in all subsequent experiments on ROAM, we choose v_0 as the neighbour of $v^\dagger$ with the *most* connections, and we connect v_0 to the $b-1$ neighbours of $v^\dagger$ with the *least* connections.

.

Figure 13: Comparing different settings of ROAM on 50 randomly generated scale-free network consisting of 100 nodes, with 3 edges added in each step of the generation process. For each such network, ROAM is executed multiple, consecutive times (the x-axis represents the number of executions). The subfigures show the source node's ranking (according to different centrality measures), and the relative change in its influence value (according to different influence models).

In the main article, due to space constraints, we only specified how the two main steps of ROAM can be applied on *undirected* networks. Next, we specify how these steps are modified to work on *directed* networks. First of all, v_0 is not chosen among the *neighbours* of $v^\dagger$, but rather among the *successors* of $v^\dagger$. This is mainly because removing a successor of $v^\dagger$ reduces its closeness centrality, whereas removing a predecessor has no such impact. As for the $b-1$ neighbours of $v^\dagger$ to be connected to v_0, they are chosen among the *predecessors* of $v^\dagger$; for each such predecessor, v_i, we add the edge (v_i, v_0). This is mainly because it could potentially rebuild the influence of $v^\dagger$ on v_0, which was hampered by the removal of the edge $(v^\dagger, v_0)$. Furthermore, for every shortest path that contains the edge $(v^\dagger, v_i)$, the addition of (v_i, v_0) could create a new alternative shortest path that does not pass through $v^\dagger$, thus further reducing the betweenness centrality of $v^\dagger$.

D.2 Experimental Results

In the main article, we only presented some of the our experimental results due to space constraints; in this subsection, we present all of our experimental results. Although most of the experimental details can be found in the main article, we add here the only missing detail, which concerns the anonymized fragments of the social networks of Facebook, Twitter and Google+ (note that the fragments of Twitter and Google+ are the only directed networks in our experiments; the remaining networks are all undirected). all anonymized fragments were taken from SNAP—the Stanford Network Analysis Platform [21].

- Facebook: the small fragment consists of 61 nodes and 272 edges; the medium one consists of 333 nodes and 2523 edges; the large one consists of 786 nodes and 14027 edges;

- Twitter: the small fragment consists of 201 nodes and 2503 edges; the medium one consists of 247 nodes and 8041 edges; the large one consists of 235 nodes and 15957 edges;

- Google+: the small fragment consists of 108 nodes and 2884 edges; the medium one contains 215 nodes and 7132 edges; the large one consists of 338 nodes and 12341 edges.

Our experimental results for the ROAM heuristic are all presented in figures 16, 17 and 18, which can be found at the end of this document.

E Constructing a Network from Scratch

Having studied the problem of disguising a node by rewiring an existing network, we now study the same problem but from a different perspective, where the goal is to construct a network from scratch, designed for the sole purpose of concealing the source node, $v^\dagger$. In this section, we will restrict our attention to undirected networks. Specifically, given n nodes, our goal is to identify a topology in which $v^\dagger$ has a reasonably-high influence, while at the same time ensuring that a certain number of nodes is ranked higher than $v^\dagger$ according to each of the three centrality measures. To tackle this problem, we propose what we call a *Lieutenant network*, the structure of which is detailed in the box below.

The Lieutenant network of size n:

- Label one node as the source node, $v^\dagger$;

- Label two groups of *lieutenants*, containing k nodes each, namely: $L = \{l_1, \ldots, l_k\}$ and $L' = \{l'_1, \ldots, l'_k\}$;

- Label all remaining nodes as $M = \{m_1, \ldots, m_\lambda\}$ where $\lambda = n - 2k - 1$; these are called *members*;

- Connect the source node to every lieutenant;

- Connect every lieutenant in L to every one in L';

- Connect every member to exactly c lieutenants from L and exactly c lieutenants from L' while ensuring that the degrees of lieutenants differ by at most 1.

Here is how the Lieutenant network works. The source node $v^\dagger$ only comes into contact with its lieutenants. These are the ones that are supposed to conceal $v^\dagger$ by ensuring that they are each ranked higher than $v^\dagger$ according to the three standard centrality measures. These are also the nodes that are supposed to pass on the influence of $v^\dagger$ to the rest of the network. Figure 14 illustrates a sample Lieutenant network with $c = 2$.
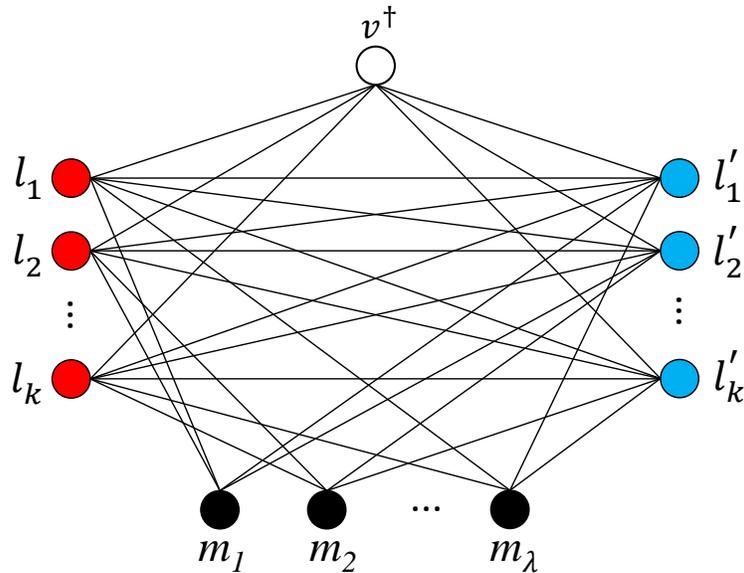

Figure 14: An illustration of a Lieutenant network with $c = 2$.

Theorem 5 *Let f denote the minimum number of members connected to any single lieutenant, i.e., $f = \left\lfloor \frac{c\lambda}{k} \right\rfloor$. Then, for every Lieutenant network such that $f > k - 1$ and $f^2 > 4ck$, all lieutenants have greater degree, closeness and betweenness centrality than the source node $v^\dagger$.*

Proof. Starting with degree centrality, the degree of the source node, $v^\dagger$, is $c_{degr}(G, v^\dagger) = \frac{2k}{n-1}$, since it is only connected to lieutenants. On the other hand, the degree of a lieutenant, l_i, is $c_{degr}(G, l_i) \geq \frac{1+k+f}{n-1}$, since it is connected to the source node, to all lieutenants from the other group, and to at least f members. As such, we have:

$$c_{degr}(G, l_i) - c_{degr}(G, v^\dagger) \geq \frac{f - k + 1}{n - 1}$$

Therefore, $c_{degr}(G, l_i) > c_{degr}(G, v^\dagger)$ for all $l_i \in L \cup L'$ when $f > k - 1$.

Moving on to closeness centrality, for any given node, v, this centrality depends inversely on the sum of the lengths of shortest paths from v to every other nodes, i.e., $\sum_{u \in V} d_G(v, u)$. For both the source node and every lieutenant, the distance to every other node is either 1 or 2. More precisely, for every $v \in \{v^\dagger\} \cup L \cup L'$, we have: $\sum_{u \in V} d_G(v, u) = 1|N(v)| + 2(n - |N(v)|) = 2n - |N(v)|$. Consequently, whenever all lieutenants have greater degree centrality than $v^\dagger$, they must also have greater closeness centrality than $v^\dagger$. This in turn implies that $c_{clos}(G, l_i) > c_{clos}(G, v^\dagger)$ for all $l_i \in L \cup L'$ when $f > k - 1$.

Finally, regarding betweenness centrality, let $\delta(v)$ denote: $\sum_{u,w \in V\setminus\{v\}:u \neq w} \frac{|\{p \in sp_G(u,w):v \in p\}|}{|sp_G(u,w)|}$. Then the betweenness centrality of a node $v \in V$ can be written as: $c_{betw}(G, v) = \frac{2}{(n-1)(n-2)}\delta(v)$. Furthermore, for any two lieutenants, $u, w \in L \cup L' : u \neq w$, let $\gamma_{u,w}$ denote the number of members that are neighbours to both of them, i.e., $\gamma_{u,w} = |M \cap N_G(u) \cap N_G(w)|$. Note that, for every pair of lieutenants belonging to the same group, the source node belongs to exactly one of the shortest path between those two lieutenants. Based on this, we have:

$$\delta(v^\dagger) = \sum_{u,w \in L:u \neq w} \frac{1}{k + 1 + \gamma_{u,w}} + \sum_{u,w \in L':u \neq w} \frac{1}{k + 1 + \gamma_{u,w}}$$

By observing that for any $a, b > 0$ we have $\frac{1}{a+b} < \frac{1}{a}$, we conclude that:

$$\delta(v^\dagger) < \sum_{u,w \in L:u \neq w} \frac{1}{k + 1} + \sum_{u,w \in L':u \neq w} \frac{1}{k + 1}$$

Now since the number of pairs of different lieutenants from each group is $\frac{k(k-1)}{2} < \frac{k(k+1)}{2}$, then:

$$\delta(v^\dagger) < k$$

Having analyzed $\delta(v^\dagger)$, let us now analyze $\delta(l_i)$ for some lieutenant $l_i \in L$ (the same analysis can be done for a lieutenant $l_j \in L'$). In particular, since l_i belongs to shortest paths (i) between every pair of lieutenants from the other group, (ii) between the source node and every member connected to l_i, and (iii) between every pair of members connected to l_i, we have:

$$\delta(l_i) = \sum_{u,w \in L':u \neq w} \frac{1}{k + 1 + \gamma_{u,w}} + \sum_{v \in M \cap N(l_i)} \frac{1}{2c} + \sum_{u,w \in M \cap N(l_i):u \neq w} \frac{1}{|sp_G(u,w)|}$$

By omitting the first term of the right-hand side of the equation and observing that $|M \cap N(l_i)| \geq f$ and $|\{\{u, w\} \subseteq M \cap N(l_i)\}| \geq \frac{f(f-1)}{2}$ and $|sp_G(u, w)| \leq 2c$ for every $u, w \in M \cap N(l_i) : u \neq w$, we end up with the following:

$$\delta(l_i) > \frac{f}{2c} + \frac{f(f - 1)}{4c} > \frac{f^2}{4c}$$

Finally, by comparing $\delta(v^\dagger)$ with $\delta(l_i)$, we find that:

$$\delta(l_i) - \delta(v^\dagger) > \frac{f^2}{4c} - k$$

Therefore, if $f^2 > 4ck$ every lieutenant has higher betweenness centrality than the source node. $\qquad \square$

As stated in the theorem, a Lieutenant network can indeed conceal its source node as far as centrality is concerned. On the other hand, as far as influence is concerned, we evaluate the network empirically to see how the different parameters affect the influence of the source node. To this end, given a Lieutenant network of 400 nodes, we varied the parameters of the network, namely k (the size of each lieutenant group) and c (the number of lieutenants from each group, connected to any given member). For every pair, (k, c), we measured the difference in centrality between the source node, $v^\dagger$, and any given lieutenant (the greater the difference, the more $v^\dagger$ is disguised), and measured the influence of $v^\dagger$ to see how this influence is affected by the disguising process.

The results are depicted in Figure 15, where the x-axis represents k and the y-axis represents c.[4] Roughly speaking, the results can be categorized into four categories:

- *small k and small c*: This yields relatively high levels of disguise in terms of betweenness, but not in terms of degree and closeness. On the other hand, it yields rather low levels of Independent-Cascade influence and Linear-Threshold influence;

- *small k and large c*: This yields relatively high levels of disguise in terms of degree and closeness, but not in terms of betweenness. On the other hand, it yields relatively high levels of Independent-Cascade influence, but not Linear-Threshold influence;

- *large k and small c*: This yields relatively low levels of disguise in terms of degree, closeness and betweenness. On the other hand, it yields relatively high levels of Linear-Threshold influence, but not Independent-Cascade influence;

- *large k and large c*: This yields relatively low levels of disguise in terms of degree, closeness and betweenness. On the other hand, it yields relatively high levels of Independent-Cascade influence, but not Linear-Threshold influence.

For future work, it would be interesting to identify other network structures that manage to disguise the source node according to all three centrality measures, while at the same time maintaining high levels of influence according to both models of influence.
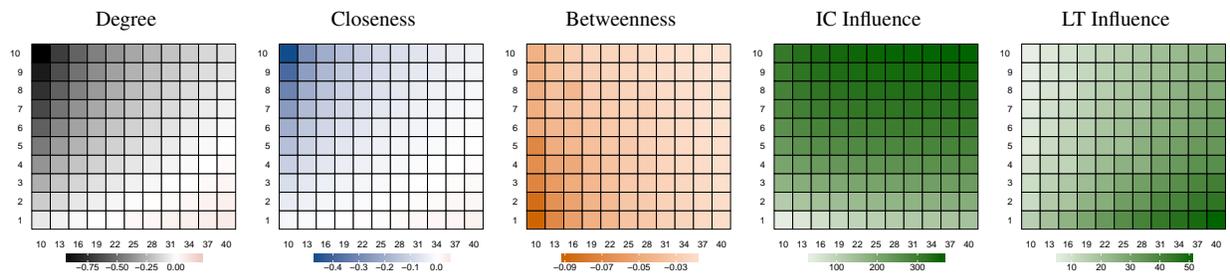

Figure 15: Given a Lieutenant netowrk of 400 nodes, with different values of parameter k (the x-axis) and parameter c (the y-axis), the figure depicts the difference in centrality between the $v^\dagger$ and a lieutenant, as well as the influence value of $v^\dagger$.

---

[4]we set $max(c) \leq min(k)$ to ensure that we have a value in every cell of the grid; otherwise some cells would correspond to networks in which the are no sufficient lieutenants to connect to.
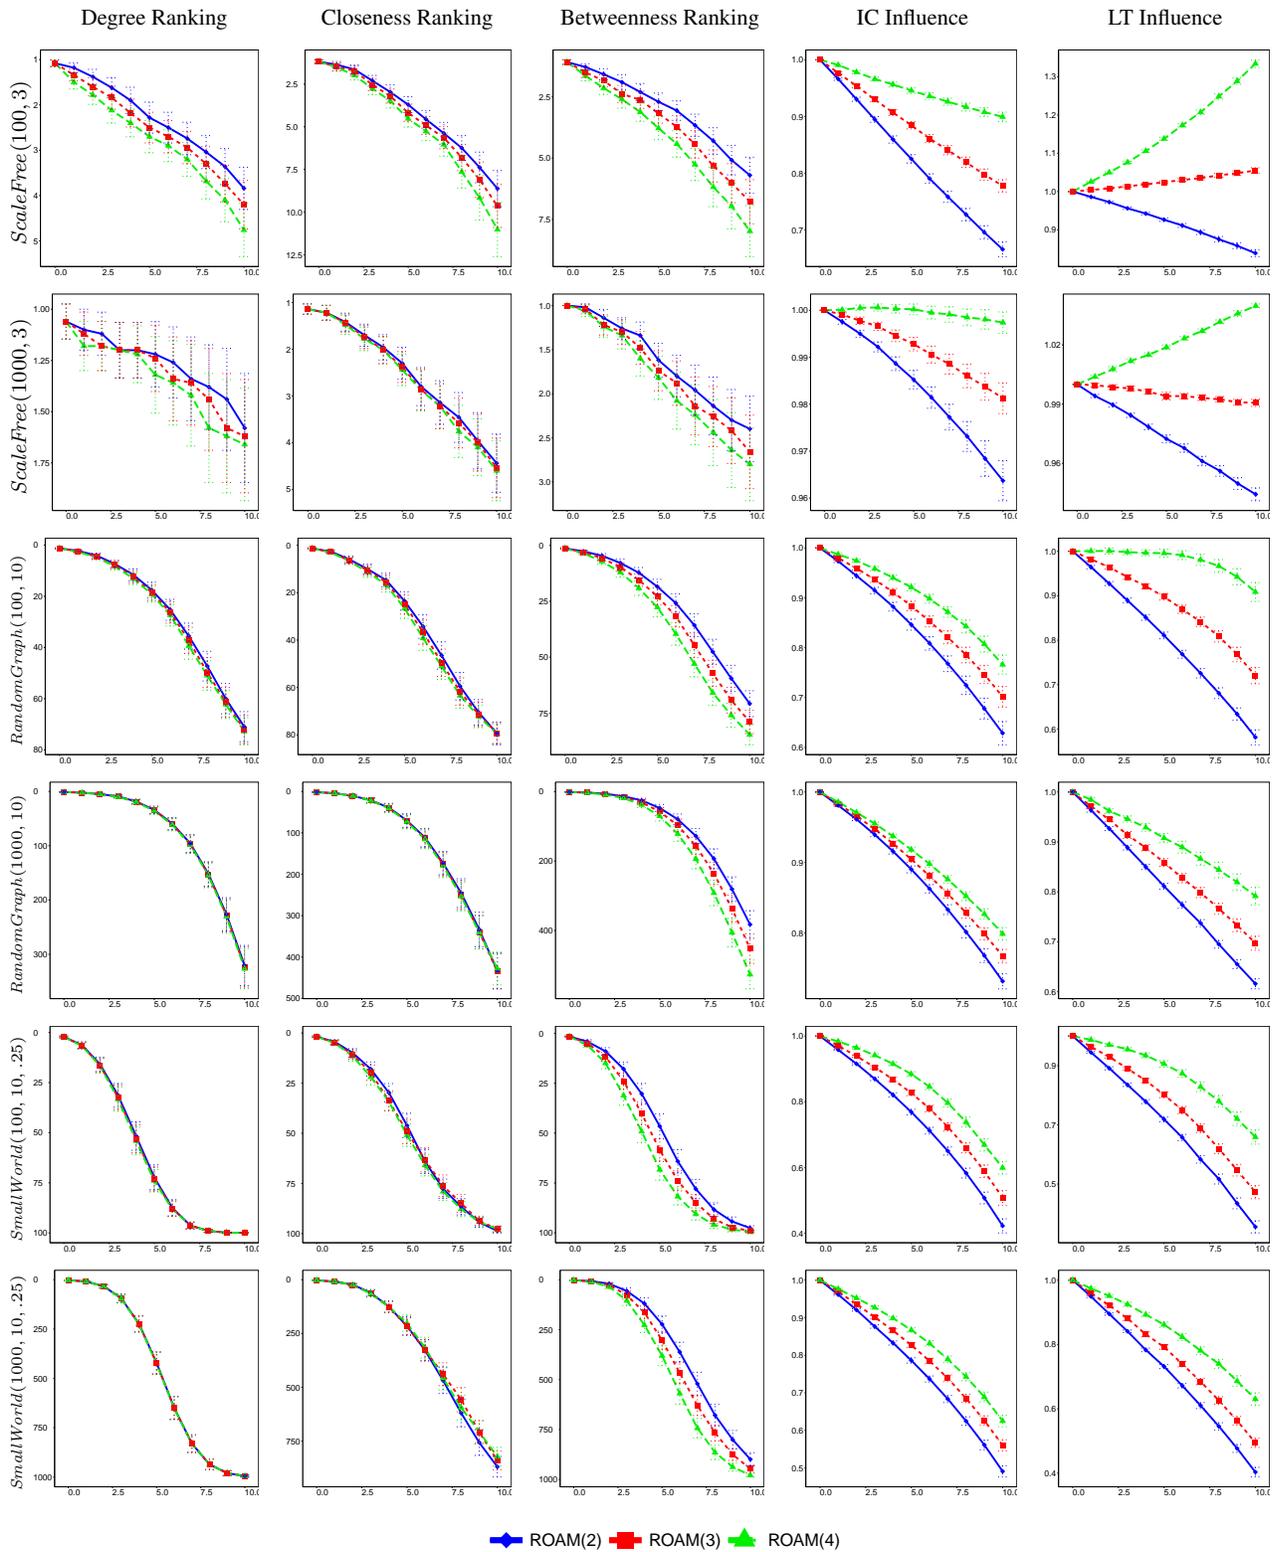
Figure 16: Consecutive execution of ROAM (the x-axis represents the number of executions). Specifically, given different random networks, the subfigures show the source node's ranking (according to the centrality measures), and the relative change in its influence value (according to the influence models). Results are shown for ROAM(b) : $b = 2, 3, 4$, where b is the budget in each execution.
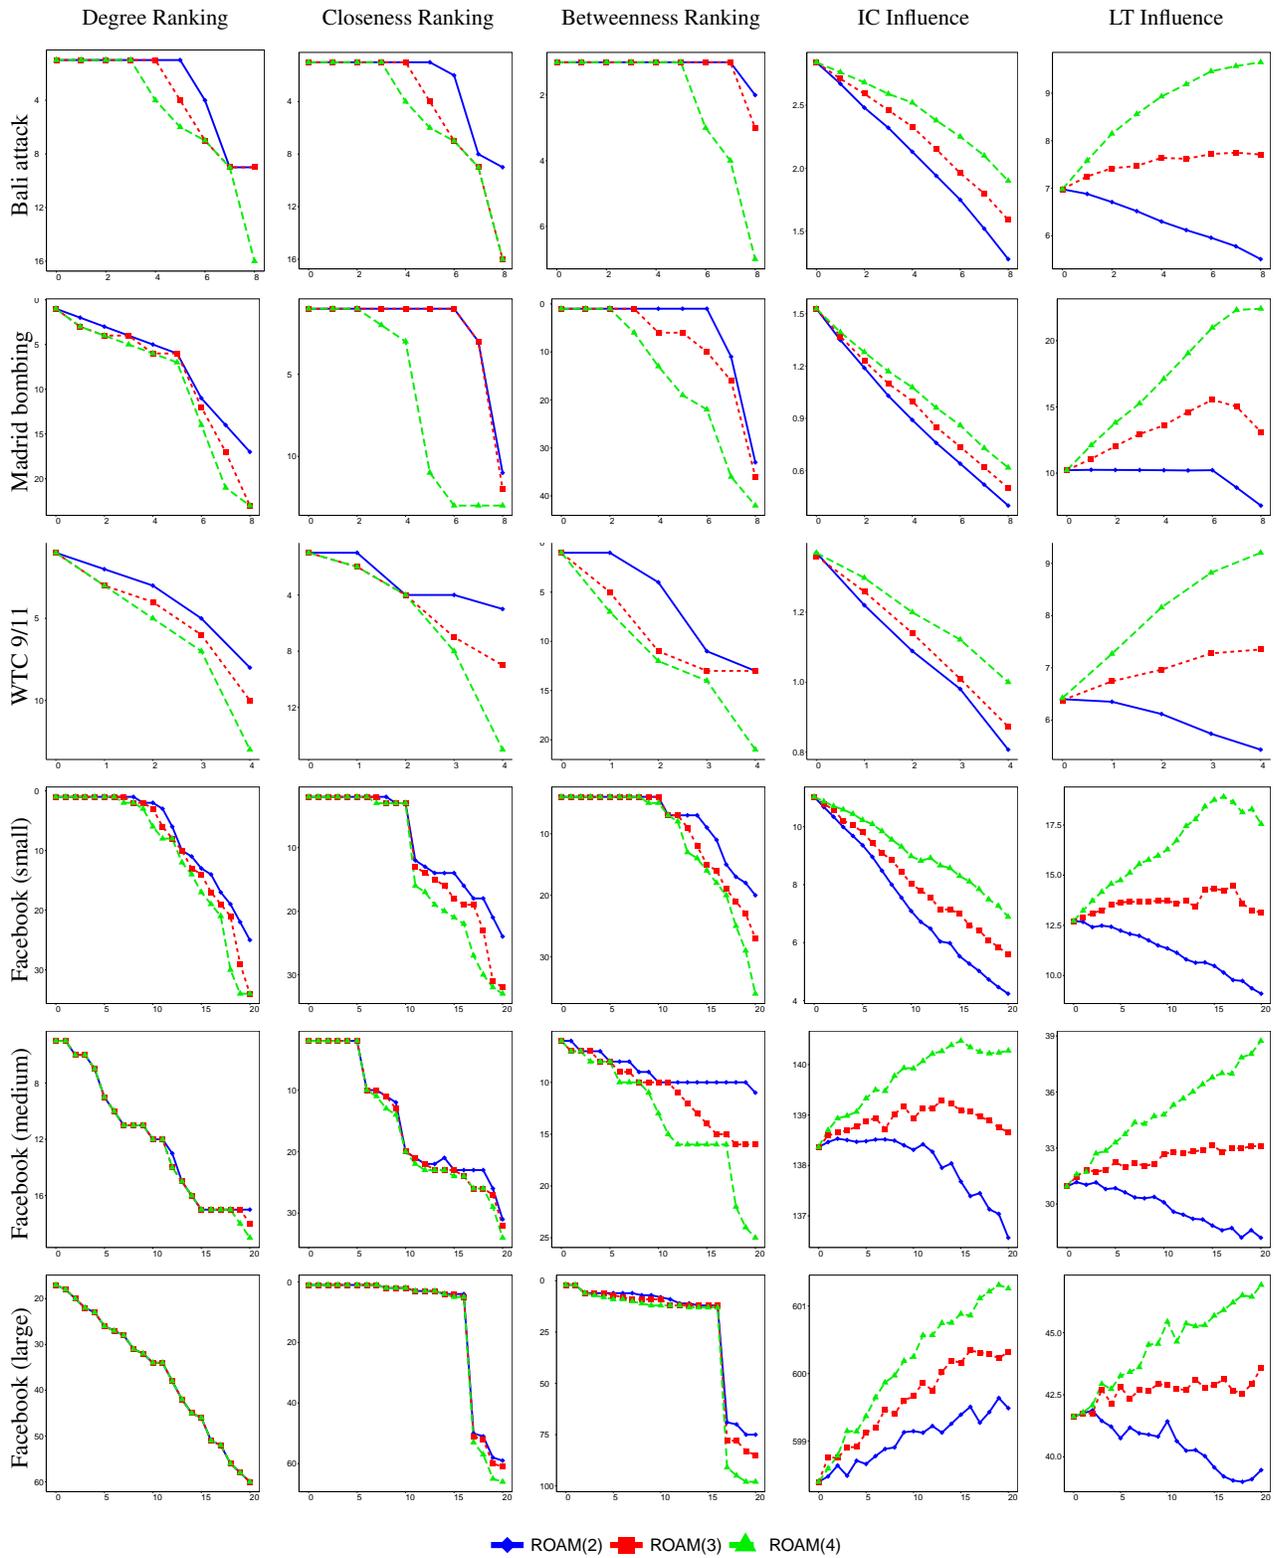
Figure 17: Consecutive execution of ROAM (the x-axis represents the number of executions). Given three terrorist networks, and different fragments of Facebook's network, the subfigures show the source node's ranking (according to the centrality measures), and the relative change in its influence value (according to the influence models). Results are for ROAM(b) : $b = 2, 3, 4$, where b is the budget in each execution.
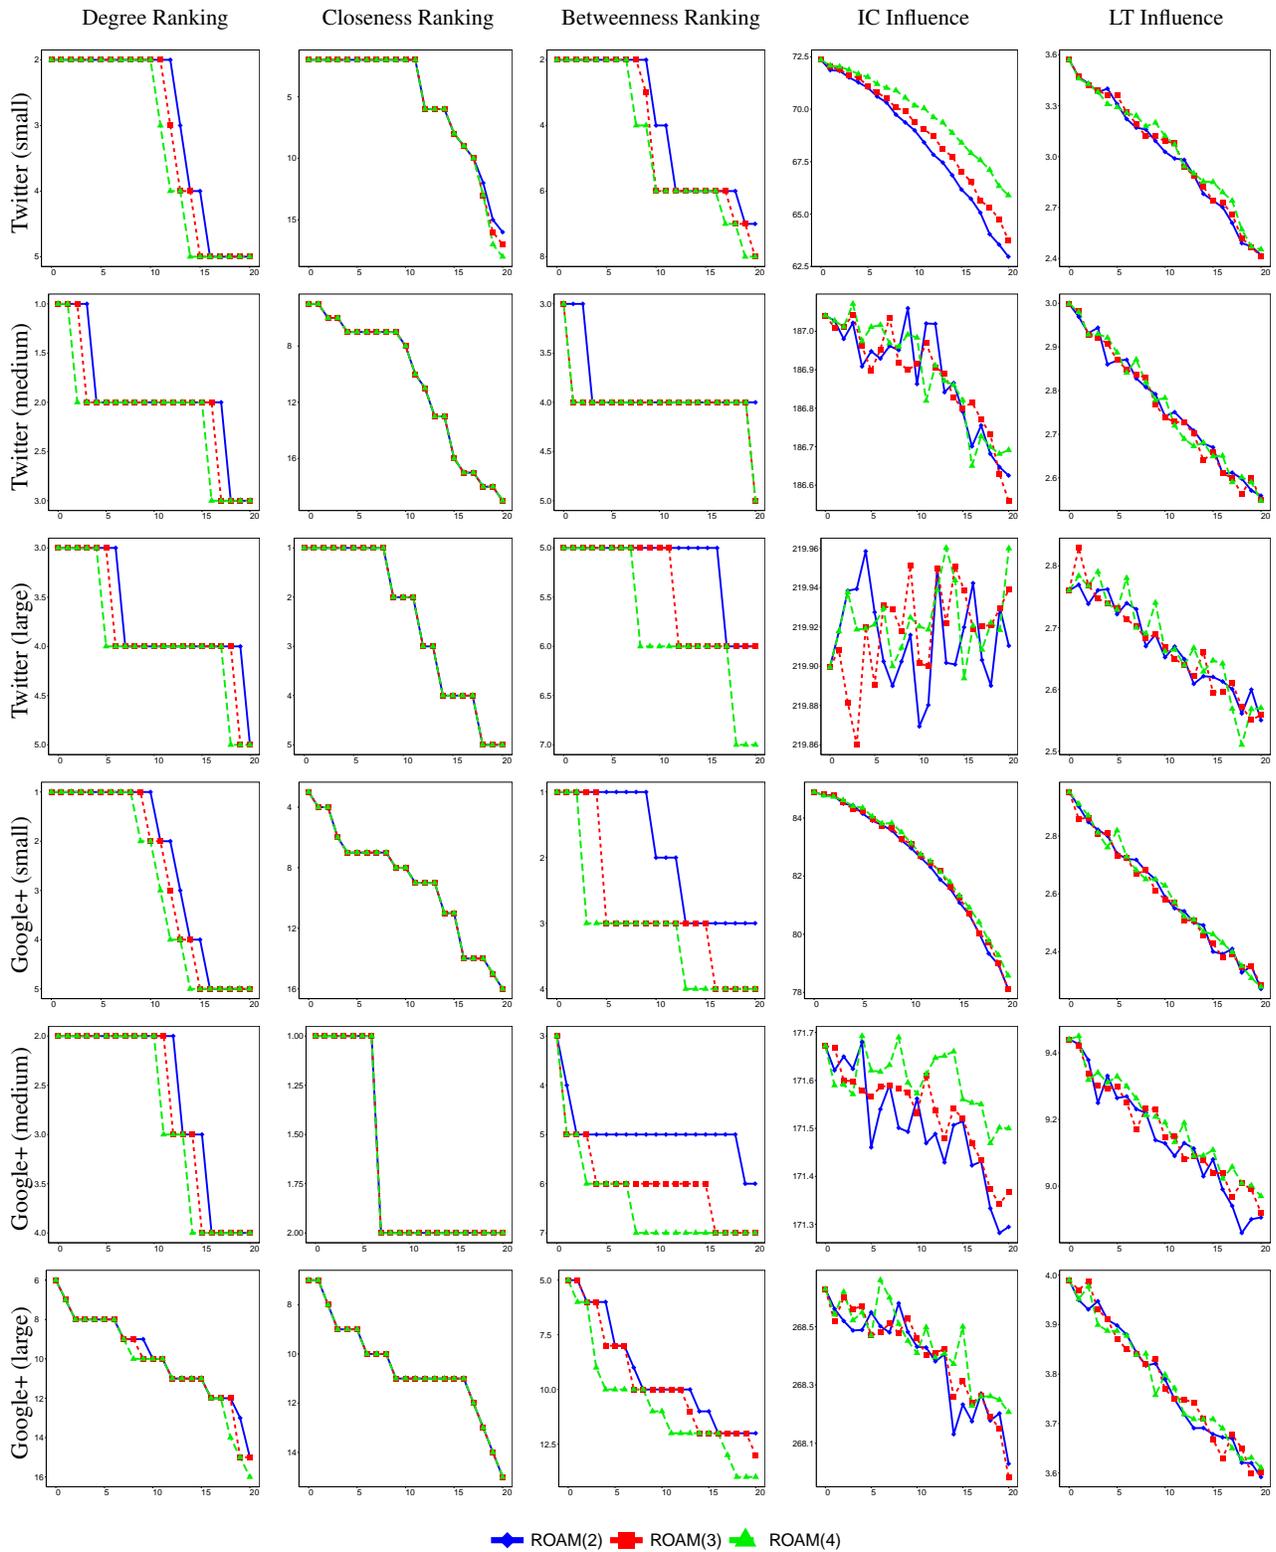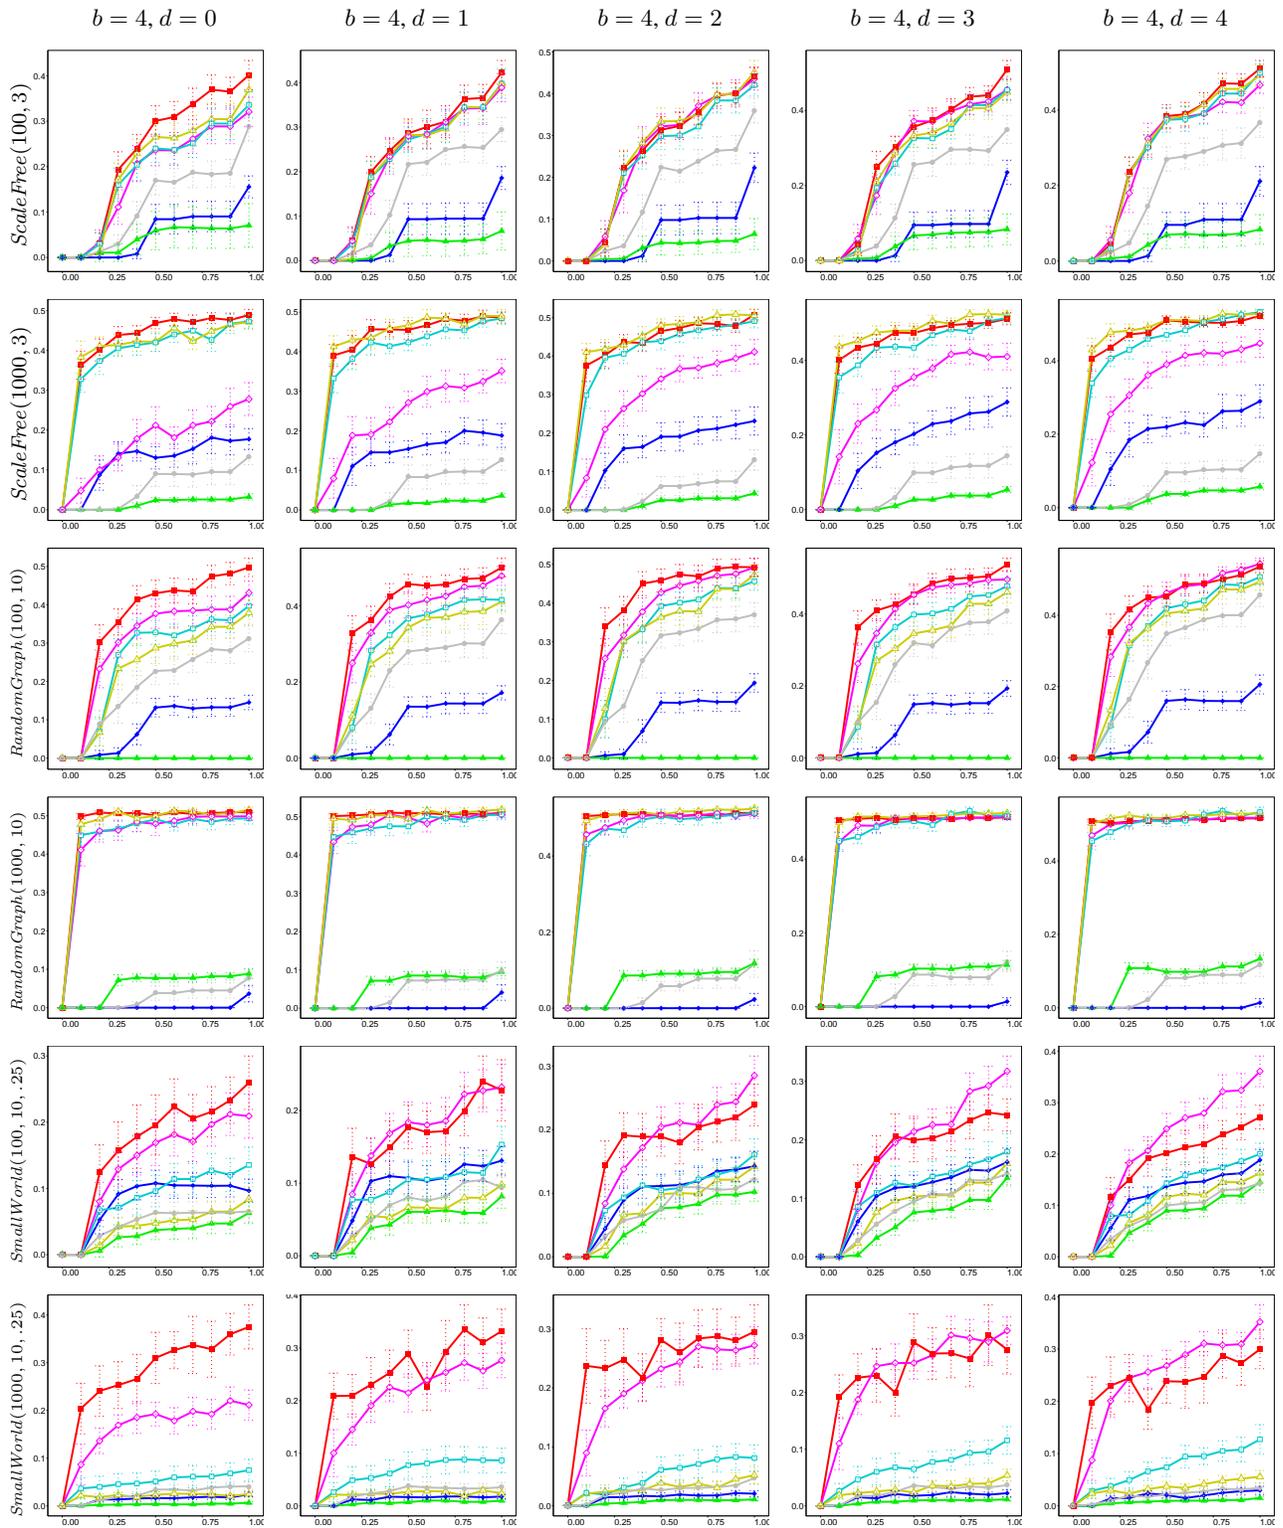
Figure 18: Consecutive execution of ROAM (the x-axis represents the number of executions). Given different fragments of the social networks of Twitter and Google+, the subfigures show the source node's ranking (according to the centrality measures), and the relative change in its influence value (according to the influence models). Results are for ROAM(b) : $b = 2, 3, 4$, where b is the budget in each execution.

Figure 19: Executing DICE multiple, consecutive rounds (the x-axis represents the percentage of completed rounds) in undirected random networks.
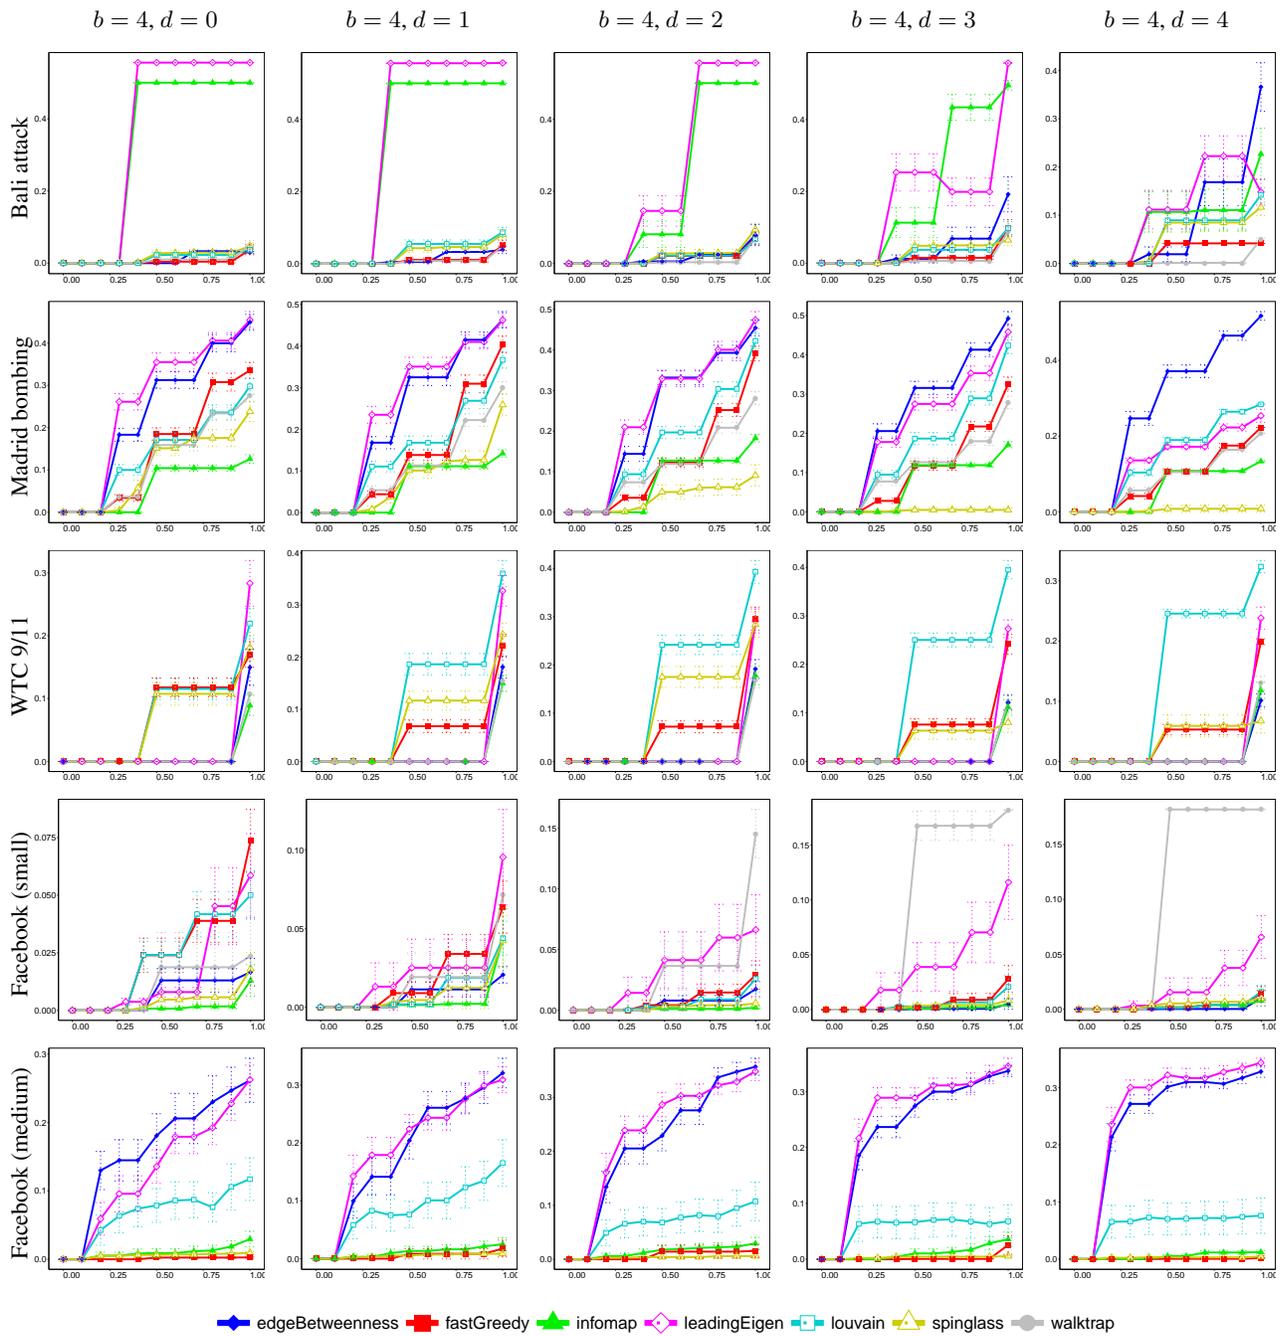
Figure 20: Executing DICE multiple, consecutive rounds (the x-axis represents the percentage of completed rounds) in undirected real-life networks.
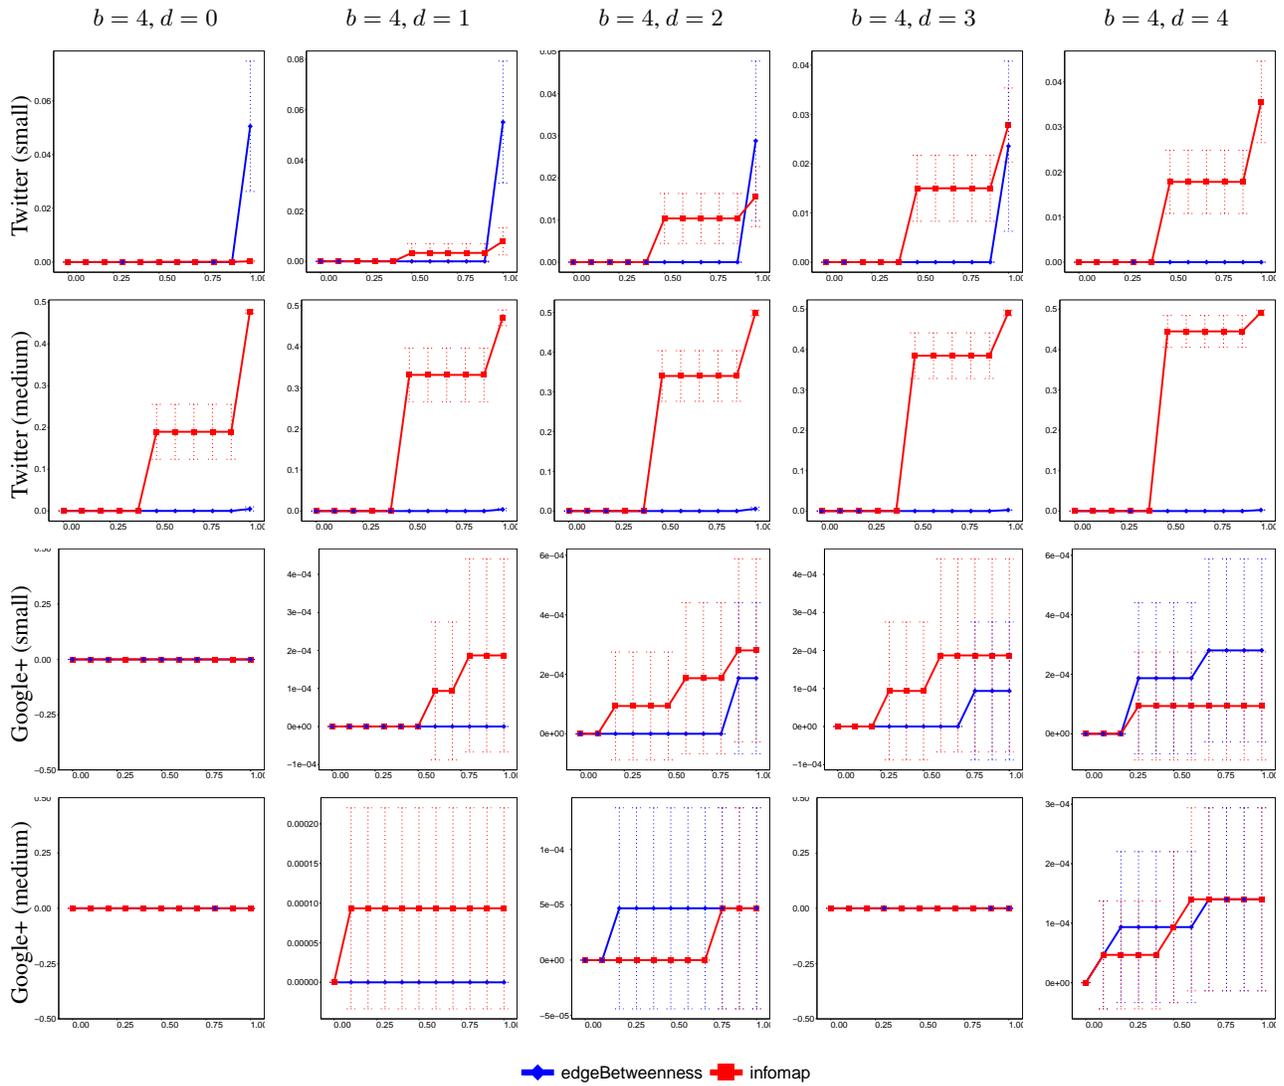
Figure 21: Executing DICE multiple, consecutive rounds (the x-axis represents the percentage of completed rounds) in directed real-life networks.
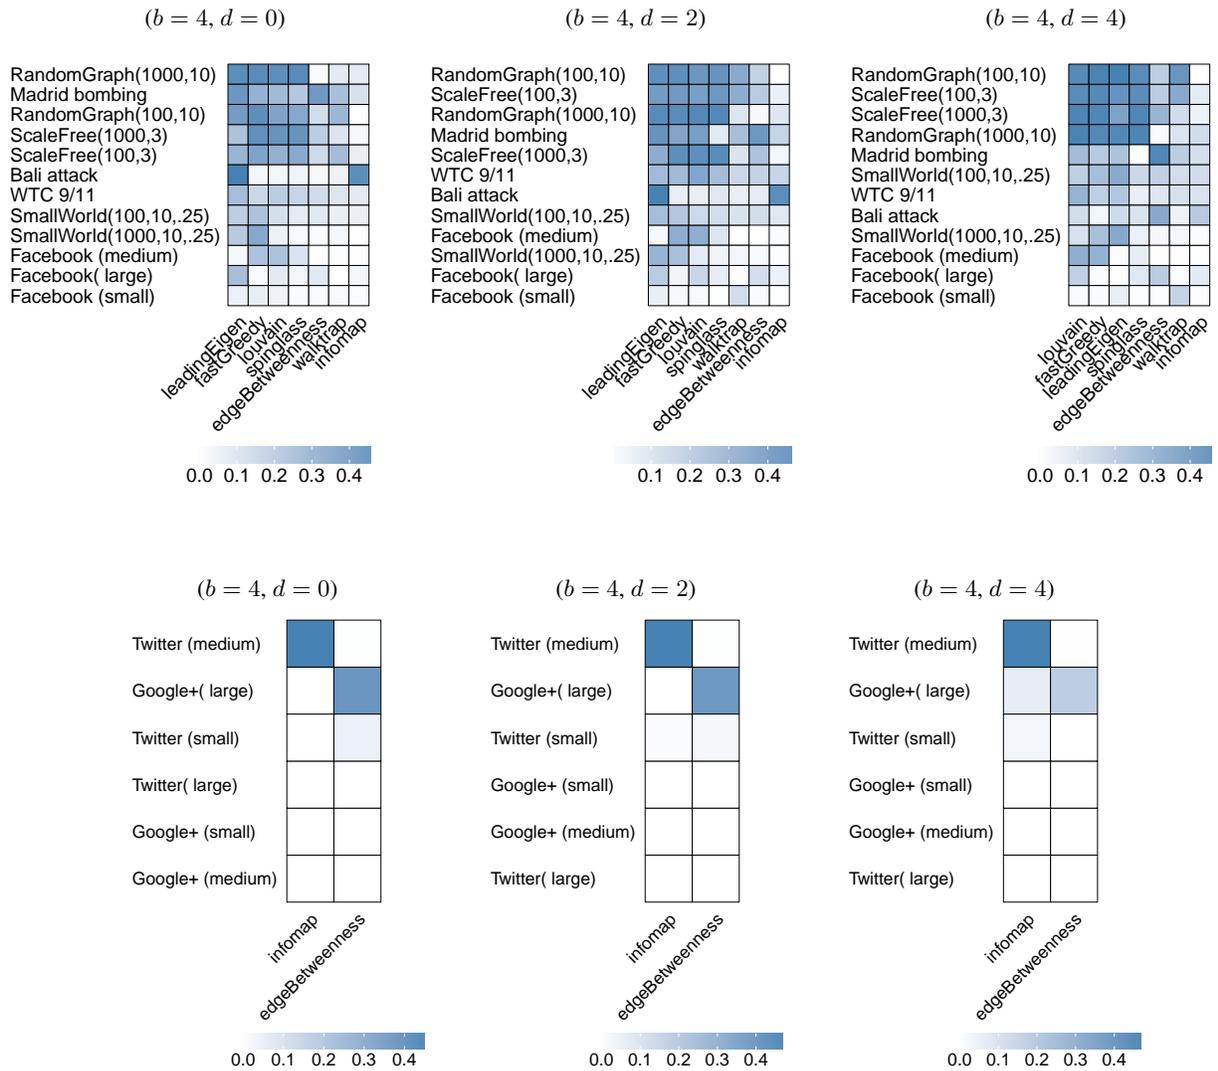
Figure 22: Average concealment-measure value in each experiment. The results for the directed networks (namely the fragments of Twitter and Google+) are presented separately.