

Introducing innovation in social networks: a cost-benefit analysis of entry point selection

Abstract—Social networks have been growing and evolving from mere means of communication into the biggest potential global market and access platform to hundreds of millions of customers ever built. However, although companies and organisations can have access to millions of potential customers almost in an instant, being able to identify the best initial entry points for introducing innovation (be it a service or product) is key to aiding its acceptance and enhancing its prospects of further diffusion into the market. In this paper, by using the economic model of return to scale, we investigate a mechanism for identifying these potential best initial entry points for introducing innovation in social networks in terms of its efficiency and a cost-benefits analysis. We present a set of experiments based on two real social network datasets and also a synthetic one that shows the effects of deploying our mechanism.

I. INTRODUCTION AND BACKGROUND

Social networks have rapidly grown to become the largest potential market ever to have emerged. Hundreds of millions of people worldwide are seen as potential customers by firms. However, although firms can have access to growing numbers of customers through social networks, new challenges have emerged as being able to identify the best initial entry points for introducing innovation (be it a service or product) is a key issue. Traditional approaches to introducing innovation are no longer applicable, since the nature and dynamics of such networks require more sophisticated solutions.

Although there exist different market techniques and models that help shed light on the process of introducing innovation, the majority of them are mostly ad-hoc solutions that may apply only to particular types of products or specific markets. Most opinions on this topic agree on the two main characteristics that most significantly affect an innovation's success: *adoption* and *diffusion*. Seminal works by [Bass1969] and [Rogers1995] have progressed on the field, but they are oriented towards traditional markets. Moreover, the focus of these works is on what happens once innovation is introduced into the market, but not before, since they do not examine what could be the best entry points to initiate product adoption and diffusion or how a firm would be able to identify these in the first place and/or choose among them. So although there is increasing interest in social networks, work on innovation seeding in such networks has been limited.

The linear threshold model proposed by [Granovetter1978] postulates the existence of a subset of individuals in a network – the *seed set* – who have already adopted the innovation. The other members of the social network subsequently adopt the innovation if the fraction of their neighbours that have already adopted it is above a certain threshold (which may differ for individuals). However, how this initial seed set is identified is

not determined. One trend in research on diffusion is to investigate when innovations become persistent in the population. For example, [López-Pintado2008] presents a method to find the threshold for the spreading rate above which a behaviour spreads and becomes persistent in a certain population. This work suggests that this threshold depends on the connectivity distribution of the social network. Similarly, [Valente1996] is focused on threshold models of collective behaviour which explain how users can eventually have different rates of adoption. Accordingly, there exist two levels of innovation rates for a user: one *macro*, relative to the system, and one *micro*, relative to her personal network. In [Kempe *et al.*2003], [Bakshy *et al.*2011], the authors conclude that there exist some influential nodes that foster diffusion of innovation throughout the network. They present different probabilistic models that allow estimating the node activation from the signals received by other users. [Aral and Walker2012] show that influential individuals tend not to be susceptible to influence from others, while susceptible individuals tend not to be influential. Another relevant approach is that of [Deroian2002], in which the author explains how the formation of the network affects the diffusion process of the innovations.

Other works have studied the effect of clustering on innovation diffusion such as on complex contagions (where individuals have high thresholds) and epidemics (where individuals have low thresholds) [Centola *et al.*2007], [Centola2010]. [Acemoglu *et al.*2011] introduce an extension to the notion of cohesiveness and characterise the final set of adopters in terms of what they call cohesive subsets. In their experimental work, they assume the individuals in the initial seed set are uniformly distributed over the network. They then show that highly clustered networks may not necessarily be more advantageous over less structured networks that contain large numbers of random links. If there exists an initial seed node within a cluster, this promotes diffusion, but if there is no such initial seed node, then it is not as easy to penetrate the cluster. In [Seeman and Singer2013], the authors present an algorithm which provides a constant factor approximation to the optimal adaptive policy for any influence function in the triggering model. In [Luu *et al.*2012], a probabilistic model for the diffusion process is presented and the conclusion reached is that the degree distribution may dynamically change.

Our work is different in that we focus on identifying the initial seed set: the most appropriate entry points to initiate the innovation adoption and diffusion. To this end, we mine the underlying structure of the network and utilise characteristics of individual nodes such as influencing power to determine the most promising entry points for seeding the innovation. The aim of this paper is two-fold: i) to provide a mechanism

for identifying suitable entry points for seeding innovation in social networks in order to aid *its early adoption and diffusion*; and, ii) to analyse the efficiency of the mechanism in terms of cost and benefits. We present a set of experiments based on two real social network datasets and a synthetic one that shows the effects of deploying the mechanism.

II. NETWORK AND INDIVIDUAL MODELS

This section presents a model for the social network, as well as a model for the innovation adoption of the user.

A. Social network model

A social network Ω can be modelled as a directed graph $G_\Omega = (V, E)$, in which V corresponds to a set of users and E is a set of edges representing relationships in Ω . Let $(v_i, v_j) \in E$, with $v_i, v_j \in V$, an edge in G_Ω representing a relationship from v_i to v_j . These relationships might denote, for example, follower-followee relations (e.g. Twitter), friendship (e.g. Facebook), job links (e.g. LinkedIn), and so on. The neighbourhood of a user is denoted as a function $neigh(v_i) = \{v_j\}$ iff. $\forall v_j \exists (v_i, v_j) \in E$.

We assume that users can communicate if there exists a link between them, hence, if user v_i wants to communicate with v_j , then there must exist a link $(v_i, v_j) \in E$. Users communicate via *signals*: a signal sent by user v_i - right after adopting a certain innovation k - to another user v_j is denoted by $\sigma_{v_i \rightarrow v_j}^k$. We assume a signal σ^k is broadcasted to all of a user's neighbours when adopting the innovation k . An example of this type of communication is Twitter: when users tweet messages, these are broadcasted to all their followers.

There are three types of stakeholders in our model. **Users** (S): are typical agents in a social network that may or may not adopt an innovation and, if so, communicate it to their neighbours. **Innovation creators** (IC): are agents that create innovation (e.g. firms). **Network explorers** (NE): are agents managed and deployed by the ICs and they are responsible for exploring the network to identify promising users through which to introduce the innovation into the network.

B. User behavioural model

Users are rational entities and when a user receives a signal about an innovation, it reasons about the presented opportunity and comes to a decision to adopt it or not. To model this reasoning process of adoption, we adapt the model proposed by [Bass1969]. In the Bass Model (BM), the probability of innovation adoption of a user is calculated as follows:

$$A_k(t) = p + \frac{q}{N} \cdot N^k(t) \quad (1)$$

p is the *coefficient of innovation* (how innovative the user is), q is the *coefficient of imitation* (how much the user is influenced by others), $N^k(t)$ is the total number of adopters up to time t , and N is the total number of potential product buyers.

In our case, we regard users as only having a partial view of the network, i.e. they only have knowledge about their neighbours. Note that since we consider that different

innovations may take place at the same time, A_k denotes the probability of adopting innovation k , while N^k is the number of neighbours that have already adopted the specific innovation k .

We assume that any link in the social network entails some sort of influencing power among the users. For instance, in Twitter, the action of following someone means the followed has power over the follower, the former is somehow appealing to or influential on the latter. Besides the inherent relation given by the link between two users in the social network, profiles could be utilised to obtain similarity measures among individuals since common interests usually foster the innovation adoption. In order to include similarity between users into the model, we extend equation 1 to represent the probability of adoption. Given a received signal $\sigma_{v_i \rightarrow \cdot}^k$, the resulting function is presented in equation 2.

$$A_k(t) = \alpha \cdot p + \beta \cdot \left[\frac{q}{|neigh(\cdot)|} \cdot N_{neigh(\cdot)}^k(t) \right] + \gamma \cdot sim(\cdot, v_i) \quad (2)$$

where $N_{neigh(\cdot)}^k(t)$ is the number of neighbours that already adopted the innovation up to time t^1 . Parameter α determines the sensitivity of the user to internal forces (coefficient of innovation) versus external forces (coefficient of imitation) to adopt the innovation (parameter β). γ represents the degree of belief the user has on the importance of the similarity with its neighbours. Note that $\alpha + \beta + \gamma = 1$ is a requirement. We call this the Adapted Bass Model (A-BM). Note that this function can cover different types of social networks. While in Facebook similarity can be obtained through the users' profiles, in Twitter there is no such notion of profile features ($\gamma = 0$).

We represent user profiles as an n -dimensional preference vector $\rho^{v_i} = (\rho_1^{v_i}, \rho_2^{v_i}, \dots, \rho_n^{v_i})$ with $\rho_j^{v_i} \in \{0, 1\}$ denoting that the user v_i presents this feature (1) or it does not (0). We calculate the similarity (sim function in eq. 2) by using the Tanimoto similarity index [Tanimoto1958]. This function is typically used to compare chemical molecules that can be represented with an array of binary elements in chemical processes. Since ρ^{v_i} profiles are binary arrays, this index appears to be an efficient way to calculate similarity, but other methods could be used here depending on the complexity and representation of the profile. Let users u and v be represented by bitmaps ρ^u and ρ^v , ρ_i^u and ρ_i^v be the i -th bit of ρ^u and ρ^v , respectively. Let \wedge, \vee be bitwise *and* and *or* operators respectively. Thus the Tanimoto similarity index is given by: $sim(u, v) = \frac{\sum_i \rho_i^u \wedge \rho_i^v}{\sum_i \rho_i^u \vee \rho_i^v}$. We assume that the contact between an IC and a user to introduce a certain innovation k into the network has an associated cost for the IC given by equation 3:

$$c(k) = c_1 + c_2 \quad (3)$$

$c_1 \in \mathbb{R}$ represents the cost for advertising the innovation (send the signal to a user) and c_2 is the cost associated with the consumed resources used by the IC to persuade a user to adopt the innovation (incentive). $c_2 = 0$ represents no need for incentives. c_2 may be obtained with a function taking into account the relative importance of the user in the network.

¹The notation (\cdot) refers to the user doing the calculation.

We define an example of this type of function in Section IV when setting up a test scenario for empirical evaluation. Note that the cost is only applicable when there exists any sort of communication between the IC and a user is selected as entry point; the diffusion phase does not have any cost for the IC.

III. THE SISM MECHANISM

The proposed mechanism for identifying and choosing suitable entry points for innovation in a social network relies on the identification of well positioned users (from a structural perspective) [Goldenberg *et al.* 2009], and those that have the ability to influence others [Fasli2006] in adopting an innovation. Since diffusion depends on the users' acceptance of the innovation, we focus on the innovation seeding, which is the only process the IC has control over. We call our approach Social Innovation Seeding Mechanism (hereafter SISM).

IDENTIFYING HUBS. We use different measures from social networks based on the well-studied property of centrality to infer which are the most promising users to introduce the innovation into the network. The concept of centrality encapsulates "micro" measures that allow us to compare nodes and to say something about how a given node relates to the overall network [Jackson2008]. Many different measures of centrality have been developed, and each tends to capture different aspects of the position that a node has, which can be useful when working with information flows, infection transmission, and other behaviours in a network. We use the following centrality measures to identify hubs:

Degree: it represents the number of links that a node has. The degree centrality of a node $v_i \in V$ is: $DC^{v_i} = \frac{d(v_i)}{|V|-1}$, where $d(v_i)$ denotes the out-degree of node v_i in the network.

Betweenness: it is a measure based on how well a user is situated on the paths it lies on [Freeman1977]. Let $np(v_j, v_k)$ be the number of paths between $v_j \in V$ and $v_k \in V$ and $np_{v_i}(v_j, v_k)$ the number of paths between v_j and v_k on which $v_i \in V$ lies on. Then we obtain the centrality of node v_i in terms of connecting v_j and v_k as the ratio $np_{v_i}(v_j, v_k)/np(v_j, v_k)$. Generalising to obtain the betweenness centrality of node v_i :

$$BC^{v_i} = \sum_{v_j \neq v_k \neq v_i} \frac{np_{v_i}(v_j, v_k)/np(v_j, v_k)}{(|V|-1) \cdot (|V|-2)/2}$$

PageRank: this well-known algorithm measures the importance of a node in terms of the fraction of time spent at that node relative to all other nodes in a random walk throughout the network. Formally: $PR^{v_i} = (1 - \lambda) + \lambda \sum_{j=1}^n \frac{PR^{v_j}}{d(v_j)}$, where λ is an attenuation parameter (we set it to 0.15 in the experiments), and $d(\cdot)$ is the out-degree of a node.

The mechanism will use one of these measures in order to identify a ranking of "central" users that essentially comprise hubs in the network, having a preference for introducing the innovation through the users in the top of this ranking.

IDENTIFYING POWER. We propose a method to assess the power that a subset of selected hubs (typically the top ones in the aforementioned ranking) exercise on others. We adhere to the notion of *referent* power [Fasli2006], as power deriving from identification. We say a user v_i has power² on user v_j iff.

there exists a connection $(v_i, v_j) \in E$ and, when v_i sends a signal $\sigma_{v_i \rightarrow v_j}^k$ user v_j automatically adopts innovation k . Then, we can say user v_i somehow influences v_j to adopt k .

In order to assess a user's power, Innovation Creators (ICs) use the so-called *network explorers* (NEs). The IC will add an NE in the network connecting it to a relevant user in order to monitor and estimate the latter's power. Note that in some social networks this process may be more complex, since connecting with an individual requires bilateral agreement (e.g. LinkedIn). We assume NEs can always connect to a selected user or, at least monitor her signals. Algorithms 1 and 2 provide the steps for the power estimation. To examine

Algorithm 1 *connectExplorer*(\cdot) algorithm that introduces the explorer into the network

Require: $v_{exp} \in NE$, $v_i \in V$ the user whose power is estimated

- 1: $E \leftarrow E \cup (v_i, v_{exp})$
 - 2: **for each** v_j in $neigh(v_i)$ **do**
 - 3: $E \leftarrow E \cup (v_j, v_{exp})$
 - 4: **end for**
-

the power user v_i has in the network, we add an NE v_{exp} and connect it to v_i (Algorithm 1, line 1). Once this initial connection is made, our algorithm also connects v_{exp} to all of v_i 's neighbours (lines 2–4). Our intention with this step is for v_{exp} to be aware of the signals sent by v_i when adopting an innovation but also the signals sent by its neighbours so that we can track the power of influence on them. Algorithm 2 presents a method to calculate the power of a user v_i by taking advantage of the connections made in Algorithm 1 (line 1).

Firstly, the explorer v_{exp} listens for any signal (line 2). When v_{exp} receives a signal from v_i (line 5), this necessarily implies that its neighbours have received it as well. Therefore, each neighbour, as explained in Section II-B, must decide if it accepts the innovation signalled by v_i or not. If the innovation is accepted, a new signal will be forwarded, and so received by v_{exp} (line 7). Explorer v_{exp} will annotate different signals received from v_i 's neighbours (line 8). We then calculate the ratio of adopted innovations after v_i 's signal (line 10). The algorithm returns a value of power as a linear combination of past influence calculated for v_i and the new influence ratio calculated (line 15). Parameter δ represents sensitivity of past influence in contrast to new one. Inevitably, influence very much depends on the user's position in the network. The same user in another location may result in exercising very different influence on its neighbours. This occurs because it is the recipient that decides whether to adopt or not the innovation. In other words, the *influencer*'s power arises from the neighbours around it, which are biased to be influenced, rather than by the *influencer*'s pro-active attitude to exercise that power. Once v_{exp} has accomplished its task, it returns the influence estimate to its IC. The IC may repeat the process with other users it may consider being promising entry points for the innovation. Finally, the IC will decide on the best entry points and send a signal to the selected user(s).

²For clarity we use power to refer to referent power.

Algorithm 2 Power estimation using NEs

Require: $v_{exp} \in NE$, $v_i \in V$ the user whose power is estimated; $pastInfluence \leftarrow 0$

```

1: connectExplorer( $v_{exp}, v_i$ )
2: while true do
3:    $positive \leftarrow 0$ 
4:    $newInfluence \leftarrow 0$ 
5:   if signalReceived then
6:     for each  $j$  in  $neigh(v_i)$  do
7:       if signalReceived then
8:          $positive++$ 
9:       end if
10:     $newInfluence \leftarrow positive/|neigh(v_i)|$ 
11:  end for
12:  if  $pastInfluence == 0$  then
13:    return  $newInfluence$ 
14:  else
15:    return  $\delta \cdot pastInfluence + (1-\delta) \cdot newInfluence$ 
16:  end if
17: end if
18: end while

```

IV. EXPERIMENTAL WORK

We are interested in exploring whether SISM can help an IC identify the best entry point according to its cost-benefit. In earlier work, some of the authors explored whether SISM helps the IC identify appropriate entry points for the adoption and diffusion of innovation, and the potential impact that the underlying network structure may have on the performance of the algorithm. In this work, we explore the impact of the budget limitations and the cost to incentivise the user to adopt the innovation at the entry point.

For simplicity, we only follow through and show in action one IC trying to introduce sequentially multiple innovations. Moreover we assume an IC typically deploys one NE per hub to be studied, in our case the top 10% of ranked users returned from the hub identification phase. In the first time step, the IC uses one (but it could be a combination) of the centrality methods to identify the candidate initial seed set and subsequently estimates the power of these users to identify which users will be selected as the entry points. It then sends a signal to each selected user with the innovation. The signalled user then decides to accept (or not) the innovation. If it accepts it (as per the adoption function in Eq. 2), the user will send a signal to its neighbours (considered at the next time step). If the entry point user does not accept the innovation, then the IC will offer an incentive.

The economic model used for the incentive to be paid to the initial adopters is the return to scale. Accordingly, there is a positive return to scale where there is a positive cost-benefits' non-linear relationship [Ferguson2008]. By assuming a positive return to scale, we also assume that better connected users will require higher incentives since they are supposed to be more effective entry points. For this reason, the IC has to identify the user(s) that best fit its budget constraint. The

incentive the IC will pay is calculated as:

$$Incentive = \frac{C}{pos(.)} \quad (4)$$

where $pos(.)$ is the user's position in the ranking of hubs as calculated, and C is a constant that may vary in different domains. We set $C = 1000$ for this set of experiments.

Note that based on the return to scale model, the cost of the incentive is a non-linear positive function of the position of the initial adopter. This means that a better positioned initial adopter is more capable of showing positively crescent non-linear benefits than a worse positioned initial adopter.

The IC selects the user according to its cost (which depends on the user's position), and its budget constraint. For example, suppose the budget constraint is 500. According to Eq. 4, the cost of the better positioned user is 1000. This is above the budget constraint, and hence, this user does not receive any incentive. The cost of the next best positioned user (2nd) is 500 which is equal to the budget constraint, and hence, this user receives the incentive. Thus, the smaller the budget, the lower the position of the user in the ranking that receives the incentive. After the initial entry point and incentivisation phase, the next time step starts with users processing the received signals. This is repeated every time step until no new signals are sent out, at which point the diffusion converges.

We have used three datasets for our experimental work. The first dataset [McAuley and Leskovec2012] was collected from surveying participants using a Facebook application. This network contains 333 users and 5038 links among them denoting reciprocal friendships (all edges in the graph are bidirectional) comprising 'circles' (or 'friends lists'). This dataset also includes node features (profiles) that are used by the SISM mechanism to calculate the similarity between users. The second dataset is based on Twitter [McAuley and Leskovec2012] with 242 users, but does not include profile information. These two datasets have the characteristics of scale-free networks. In such networks, there are certain sets of vertices with a degree that greatly exceeds the average (hub nodes). We have also created an artificial network of the same size as the Facebook dataset (i.e. 333 nodes) as a poorly connected network in order to see the effects of the SISM mechanism with varying budgets. We experimented with larger sizes of poorly connected networks (up to 10,000 agents), but the results are similar as for the smaller size network, hence we report these here for comparative purposes.

For our experiments, we use three different centrality measures: Degree, Betweenness and PageRank. We run nine treatments based on the combination of centrality measures and datasets. On every combination we tested seven different budget constraints (10, 25, 50, 100, 250, 500 and 1000), on ten different initial nodes. The results presented are the average of 10 runs (for each experiment) with different random seeds to set up different initial user populations (different parameters in eq. 2). Our aim is to test the ability of higher budget constraints to generate higher number of adopters or to increase the speed of the adoption compared to lower budgets. We have also tested the performance of SISM against a Random selection

strategy (what in marketing, for instance, could be compared to a broadcasting strategy) in which entry points are selected in a random manner.

As we are interested in testing the cost-benefit to incentivise initial adopters, we measure the cost benefits of the incentive comparing three different parameters among the budget constraints: the convergence rate, the total number of adopters and the initial number of adopters.

A. Results of evolution of the numbers of adopters

In this section, we show in detail the results of the experiment that analyses the convergence rate and the total number of adopters on the diffusion of innovation with the three different datasets (Facebook, Poor and Twitter). On every treatment with an individual dataset, we show the evolution of the number of adopters among the three SISM centrality measures (Degree, Betweenness and PageRank).

1) *Facebook Dataset Results:* The results shown in Figures 1 and 2 provide insights on the convergence rate and final numbers of adopters of the innovation with the Facebook dataset. Firstly, the convergence rate is very similar between different budget levels when using the Degree and PageRank centrality measures regardless the budget constraint. Note in Figure 2 that it is not possible to identify a budget constraint level that could lead to a higher convergence rate. This means that the budget constraint level has no influence on the convergence rate.

Secondly, even though the convergence rate for different budget constraint levels are similar when using the Degree and PageRank centrality measures, these convergence rates are higher than the case of random selection. This can be viewed clearly in Figure 1 as the curve of the random selection convergence is below every curve on the SISM Degree and PageRank centrality measures regardless the budget constraint level.

Thirdly, the convergence rate is different among the various budget levels when the Betweenness centrality measure is used as part of the SISM mechanism. Differently to SISM Degree and PageRank centrality measures, Figure 2 shows that different budget constraint levels lead to a different convergence rate when using the SISM Betweenness centrality measure. Note that the curves are dispersed, with some curves below the curve of the random selection convergence rate, while other curves are above the random selection.

Fourthly, Figure 2 suggests that regardless the level of the budget constraint, the final number of adopters is very similar in the three SISM centrality measures used in this treatment. This means that this worthless to incentive the well connect initial adopters if the objective is only to reach as much user as possible. For this reason, we argue that the incentive should focus only on the number of initial adopters, not on the number of final adopters.

2) *Twitter Dataset Results:* The results shown in Figures 3 and 4 provide a series of insights regarding the convergence rate and final numbers of adopters of the innovation with the Twitter dataset. Firstly, similarly to the Facebook dataset, the convergence rate is very similar between for the SISM

Degree and PageRank centrality measures and this appears to be regardless the budget constraint. Note in Figure 3 that it is not possible to identify the budget constraint level that could lead to a higher convergence rate. This means that the budget constraint level has no influence on the convergence rate.

Secondly, the convergence rate appears to differ for the various budget levels for the SISM Betweenness centrality measure. Unlike the SISM Degree and PageRank centrality measures, Figure 4 shows that different budget constraint levels lead to a different convergence rate when the SISM Betweenness centrality measure is used.

From Figure 3 note that there is peak down for low values of Budget in the SISM Betweenness centrality measure's 3D-plot. This that could mean that this worths to increase the value to incentive well connected initial adopters for this SISM centrality measure.

Finally, Figure 4 shows that the convergence rates are higher to the random selection for all SISM centrality measures, regardless the budget constraint levels. Figure 4 also suggests that the final number of adopters is very similar for the three SISM centrality measures used in this treatment regardless the budget constraint levels used.

3) *Poor Dataset Results:* As indicated earlier, in addition to conducting experiments with data from real social networks we created an artificial poorly connected network equal in size to the Facebook network. The results shown in Figures 5 and 6 illustrate the results regarding the convergence rate and final number of adopters with this "Poor" dataset. The first observation is that the convergence rate is not similar for the SISM Degree and Betweenness centrality measures and various budget constraint levels. Interestingly note the valley in the SISM Degree and Betweenness centrality measures as shown in the 3D-plot graphs of Figure 5. The analysis of this valley is not straight forward. The point is that this is not an expected result. Perhaps, this is not possible to assure that the impact of the budget constraint's level on both SISM centrality measures is not positively correlated to the level of budget constraint.

Secondly, the convergence rate is similar for the different budget constraint levels when the SISM PageRank centrality measure is used. Unlike the SISM Degree and Betweenness centrality measures, Figure 5 shows that with the SISM Betweenness centrality measure, the convergence rate is very similar in the presence of different budget constraint levels.

This can be observed in Figure 5 where the 3D-Plot for the SISM PageRank centrality measure is smooth, unlike the 3D-Plots for the SISM Degree and Betweenness centrality measures which is rather jagged. The shape of the 3D-Plot for the SISM PageRank centrality measure is expected because this is similar to the shape of 3D-Plot for this SISM centrality measure in Figure 1 and in Figure 3. The question remains on the valley present in the 3D-Plot for the SISM Degree and Betweenness centrality measure.

Thridly, Figure 6 shows that the convergence rates of the SISM Degree and Betweenness centrality measures are close to the random selection. In fact, at the start of the diffusion process and when the SISM Degree and Betweenness centrality measures are used, the random selection shows a higher




Fig. 1. Number of adopters for the Facebook dataset vs the number to ticks and the level of budget constraint for the three SISM centrality measures(Degree, Betweenness and PageRank)

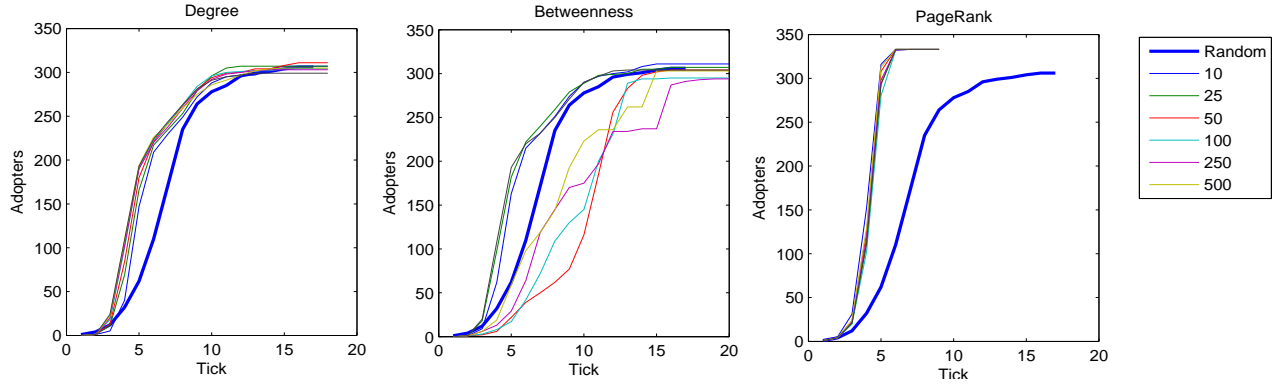


Fig. 2. Evolution of adopters on different budgets with centrality measures vs Random in the Facebook dataset




Fig. 3. Number of adopters for the Twitter dataset vs the number to ticks and the level of budget constraint for the three SISM centrality measures(Degree, Betweenness and PageRank)

number of adopters; this only changes closer to the tenth tick in time.

Finally, Figure 6 shows that the convergence rate of the PageRank centrality measure is higher to the random selection and this is irrespective of the budget constraint level used. Furthermore and in contrast to some of the results observed when the SISM Degree and Betweenness centrality measures are used, the number of adopters when using the PageRank

centrality measure and at about the tenth tick, is near 250, more than twice the number of adopters when using random selection.

B. Comparative results for initial adopters

We would like to study the relation between the budget constraint level and the number of initial adopters and in this section, we show the results of the experiment with the three

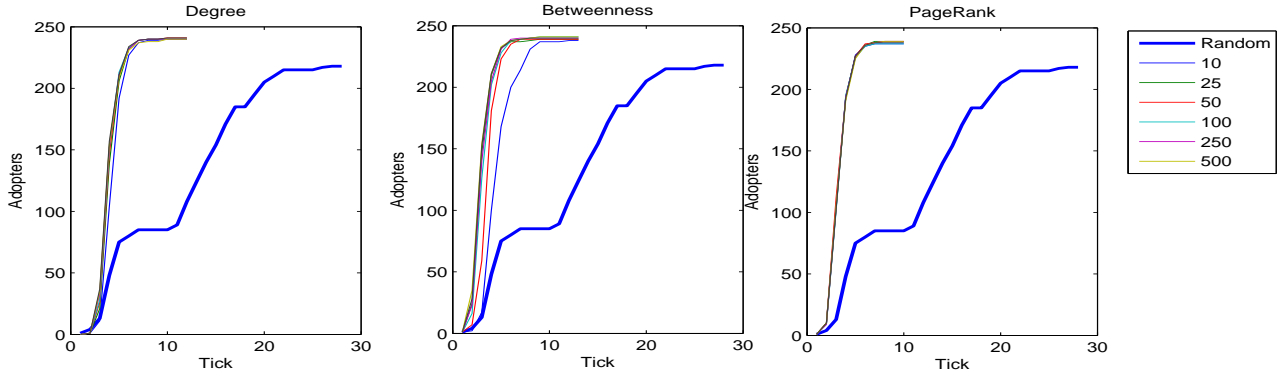


Fig. 4. Evolution of adopters on different budgets with centrality measures vs Random in the Twitter dataset




Fig. 5. Number of adopters for the Poor dataset vs the number to ticks and the level of budget constraint for the three SIS centrality measures (Degree, Betweenness and PageRank)

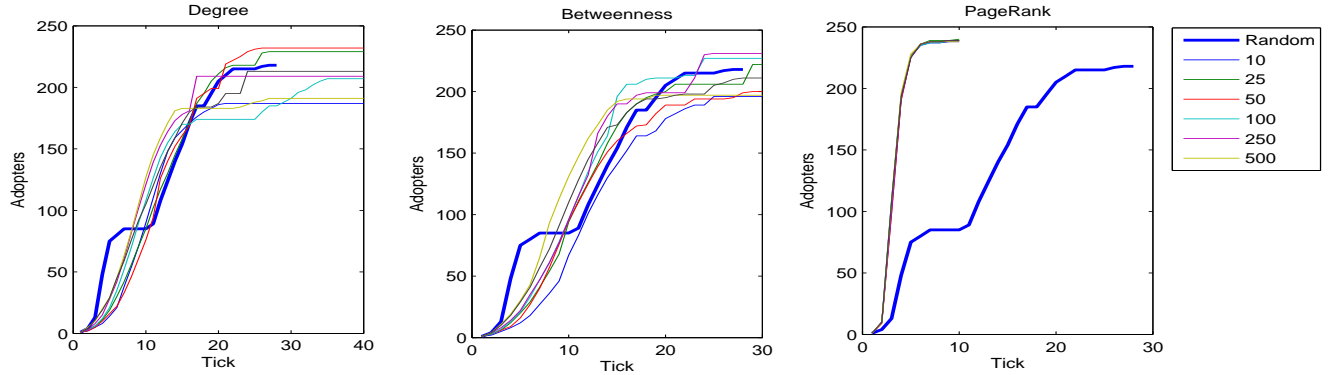


Fig. 6. Evolution of adopters on different budgets with centrality measures vs Random in the Poor dataset

different datasets (Facebook, Poor and Twitter) that analyse the impact of the budget constraint level on the number of initial adopters. In particular, we examine the number of initial adopters which comprises those adopters that after having been identified as part of the initial seed set as an outcome of one of the three centrality measures used by the SIS mechanism (Degree, Betweenness and PageRank), they were incentivised by a payment within the confines of the specific budget constraint used, and subsequent to that payment, they

adopted the innovation. We discuss the results per centrality measure used.

I think we need a few sentences here to describe how random selection was done here so that the reader can understand what we are comparing against

1) *Degree*: The results shown in Figure 7 suggest four insights that can be drawn about the diffusion of innovation using the SIS Degree centrality measure. Firstly, it appears that the budget constraint is positively correlated to the number




Fig. 7. Results for Degree centrality measure

of initial adopters. This can be seen from that the bar chart is crescent along the budget constraint on every dataset. Even though with Poor dataset, this is only evident from the budget constraint of 250 monetary unities.

Secondly, lower levels of incentives result in the number of initial adopters being inferior to the random selection entry point with the Facebook dataset. As can be observed, it is only worth incentivising some initial adopters if the budget is above 50 monetary units, otherwise, this is preferable to select randomly the initial adopters.

Thirdly, the results suggest that the use of incentives is not worthwhile with the Twitter dataset. When using random selection of entry points, the number of initial adopters is over 70, while even with the highest budget constraint, the number of initial adopters only reaches between 30 and 40.

Finally, the return to scale is decreasing with the budget constraint level. For example, while a budget of 10 monetary units results in approximately 10 initial adopters with the Twitter dataset, a budget of 1000 monetary unities (100x higher) results in approximately 30 initial adopters (3x higher). What this in essence tells us, is that depending on the underlying network structure, in some cases, it is not worth providing huge incentives (in the form of payment, free products etc.) as this will not make a significant difference in the number of initial adopters. Therefore the level of incentivisation needs to be chosen carefully.

2) *Betweenness*: Examining the results illustrated in Figure 8, we can draw the following conclusions on using the SISM mechanism with the Betweenness centrality measure. Firstly, the budget constraint is not correlated to the number of initial adopters. Unlike the discernible patterns that we can see in Figure 7 for the Degree centrality measure, the results illustrated in Figure 8 appear to have no specific pattern that can be easily deduced. For example, with the Twitter dataset, the number of initial adopters with the budget constraint of 25 monetary units is higher to the budget constraint of 50, 100 and 250 monetary units; but lower to the number of adopters with budget constraints of 500 and 1000 monetary units.

Hence, the results do not provide conclusive evidence that the use of incentives results in the number of initial adopters being higher than by using the random selected entry points.




Fig. 8. Results for Betweenness centrality measure

Indeed in the Facebook and Poor datasets the differences that we see in the number of adopters even when the budget constraint increases are very small. Counterintuitively, in the Facebook dataset a budget of 25 yields a higher number of adopters than a budget of 500.

Finally, similar to the diffusion of innovation with the SISM Degree centrality measure, the return to scale is decreasing with the budget constraint. For example, while a budget of 10 monetary units results in approximately 20 initial adopters with the Twitter dataset, a budget of 1000 (100x higher) results in approximately 160 initial adopters (8x higher).




Fig. 9. Results for PageRank centrality measure

3) *PageRank*: The results shown in Figure 9 provide two different insights about the diffusion of innovation with the SISM PageRank centrality measure. Firstly, the budget constraint is not correlated to the number of initial adopters. We can see that all of the datasets show similar number of initial adopters regardless the level of the budget constraint.

Secondly, and unlike the Betweenness centrality measure, as can be seen, with all three datasets of the experiment, the use of incentives results in the number of initial adopters being higher to randomly selected entry points. For example, even with the Poor dataset, the number of initial adopters based on incentives is around 100, while the number of initial adopters when

randomly selecting entry points is below 20 (approximately 5x lower).

Finally, similar to the two previous cases of using the Degree and Betweenness centrality measures, Figure 9 suggests that the return to scale is indifferent to the budget constraint. For example, while a budget of 10 monetary units results in approximately 20 initial adopters with the Twitter dataset, a budget of 1000 (100x higher) results in approximately 160 initial adopters (8x higher).

V. DISCUSSION

In this section, we consider the results discussed in the preceding sections and draw our conclusions. Firstly, the results suggest that there is no benefit from low budget levels because this can lead to a selection of an initial set of adopters which is inferior to the set of randomly selected adopters. This means that if a firm is interested in increasing the number of initial adopters, it will have to invest significant funds, otherwise the benefit from lower budgets is questionable. Lower budget levels entail that users that are less central will be chosen to be incentivised (as users further up the centrality ranking would require higher payments) which may further impact the diffusion process in the next step as they will inevitably be less influential.

Secondly, we can surmise that there may be a level beyond which it would be worth to incentivise some initial adopters. Even though at this stage of our work it is not possible to determine what this level may be, the results suggest that such a level may exist. This would mean that above this budget level, it would be worth offering incentives as these would increase the number of initial adopters, but below this level, a firm would be better off just selecting random entry points.

Thirdly, we cannot conclusively say whether the number of initial adopters is positively correlated to the budget constraint levels, as we do not see the same patterns emerging from the use of different budget constraint levels and centrality measures equally across all datasets. Even though we can observe a positive correlation in some combinations of datasets and SISM centrality measures, such as the Twitter dataset and the SISM Degree centrality measure, we cannot draw this as a generic conclusion.

Finally, focusing only on the positive correlated combination of dataset SISM centrality measure. The results suggest that the return to scale is decreasing to the budget constraint. For example, with the SISM Degree centrality measure, while a budget of 10 results in approximately 10 initial adopters, a budget of 1000 (100x higher) results in approximately 30 initial adopters (3x higher). It could suggest that this could be more efficient to incentive some initial adopter who has the cost that is near to the level in which it is worth to incentive adopters.

VI. CONCLUSIONS

In this paper, we investigated the problem of determining the set of users to be targeted (seed set) to initiate the processes of adoption and diffusion of an innovation. To identify the elements of the seed set we mine the underlying structure

of the network and utilise characteristics of individual nodes such as influencing (referent) power to determine the most promising entry points for seeding the innovation. In addition, we study the impact of the budget limitations and the cost to incentivise the user to adopt the innovation at the entry point. The economic model used for the incentive to be paid to the initial adopters is the return to scale which assumes that there is a positive return to scale where there is a positive cost-benefits non-linear relationship. Accordingly, better connected users will have to be paid more since they are assumed to be more effective entry points.

We have described the experiments conducted involving two real social network datasets and a synthetic one. Based on the results obtained, we can draw the following main conclusions. Firstly, the results suggest that there is no benefit from low budget levels because this can lead to an initial set of adopters selected which is inferior to the set of randomly selected adopters. Secondly, there is a budget point above which it is worth to incentive some initial adopters. Thirdly, the number of initial adopters is positively correlated to the budget constraint levels. Finally, the results suggest that the return to scale is decreasing; for example, with the SISM centrality Degree, while a budget of 10 results in approximately 10 initial adopters, a budget of 1000 (100x higher) results in approximately 30 initial adopters (3x higher).

Hence, our results have shown that using the SISM mechanism enables us to identify appropriate entry points to initiate the innovation adoption and diffusion processes. Targeting those entry points rather than random users yields better results. In particular, in relation to incentivising users to adopt the promoted innovations, we observe that this can have an effect, though there appears to be a cut-off budget level beyond which the return to scale is decreasing. Hence, budgets for incentivisation of users need to be carefully drawn.

There are several avenues for future work. Firstly, we would like to experiment with larger social network datasets. Secondly, we are interested in exploring the effects of ICs competing in the market for closely related innovations and in addition adding (product/services) preferences and budget constraints on the user side that in reality can constrain a user from adopting an innovation, no matter how attractive or desirable this innovation may be. A third direction is to study an extension of the mechanism in order to deal with dynamic social networks, in which structural properties change due to the evolution of the links in the network. Finally, a forth direction is to model the level that it starts to worth to incentive initial adopters

REFERENCES

- [Acemoglu *et al.*2011] D. Acemoglu, A. Ozdaglar, and M. E. Yildiz. Diffusion of innovations in social networks. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, pages 2329–2334, 2011.
- [Aral and Walker2012] Sinan Aral and Dylan Walker. Identifying Influential and Susceptible Members of Social Networks. *Science*, 337(6092):337–341, 2012.
- [Bakshy *et al.*2011] Eytan Bakshy, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. Everyone’s an influencer: Quantifying influence on twitter. In *Proceedings of the Fourth ACM International Conference on*

- Web Search and Data Mining*, pages 65–74, New York, NY, USA, 2011. ACM.
- [Bass1969] Frank Bass. A new product growth model for consumer durables. *Management Science*, 15, 1969.
- [Centola *et al.*2007] D. Centola, V. M. Eguiluz, and M. W. Macy. Cascade dynamics of complex propagation. *Physica A: Statistical Mechanics and its Applications*, 374(1):449–456, 2007.
- [Centola2010] D. Centola. The spread of behavior in an online social network experiment. *Science*, 329(5996):1194–1197, 2010.
- [Deroian2002] Frédéric Deroian. Formation of social networks and diffusion of innovations. *Research Policy*, 31(5):835 – 846, 2002.
- [Fasli2006] Maria Fasli. On the relationship between roles and power: preliminary report. In *Proceedings of the 2006 ACM symposium on Applied computing*, SAC ’06, pages 313–318, New York, NY, USA, 2006. ACM.
- [Ferguson2008] C. E. Ferguson. *The Neoclassical Theory of Production and Distribution*. Cambridge University Press, 2008.
- [Freeman1977] Linton C. Freeman. A Set of Measures of Centrality Based on Betweenness. *Sociometry*, 40(1):35–41, 1977.
- [Goldenberg *et al.*2009] Jacob Goldenberg, Sangman Han, Donald R. Lehmann, and Jae W. Hong. The role of hubs in the adoption process. *Journal of Marketing*, 73(2):1–13, 2009.
- [Granovetter1978] Mark Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978.
- [Jackson2008] Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, USA, 2008.
- [Kempe *et al.*2003] David Kempe, Jon Kleinberg, and Éva 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, New York, USA, 2003. ACM.
- [López-Pintado2008] Dunia López-Pintado. Diffusion in complex social networks. *Games and Economic Behavior*, 62(2):573 – 590, 2008.
- [Luu *et al.*2012] Duc Minh Luu, Ee-Peng Lim, Tuan-Anh Hoang, and Freddy Chong Tat Chua. Modeling diffusion in social networks using network properties. In *Proceedings of the 6th ICWSM*, 2012.
- [McAuley and Leskovec2012] Julian J. McAuley and Jure Leskovec. Learning to discover social circles in ego networks. In *26th Annual Conference on Neural Information Processing Systems*, pages 548–556, 2012.
- [Rogers1995] E. Rogers. *Diffusion of innovations*, chapter Elements of diffusion, pages 1–37. The Free Press, 1995.
- [Seeman and Singer2013] Lior Seeman and Yaron Singer. Adaptive seeding in social networks. *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 459–468, 2013.
- [Tanimoto1958] T. Tanimoto. An elementary mathematical theory of classification and prediction. IBM Internal Report, 1958.
- [Valente1996] Thomas W. Valente. Social network thresholds in the diffusion of innovations. *Social Networks*, 18(1):69 – 89, 1996.