

RESEARCH

When Can Overambitious Seeding Cost You?

Shankar Iyer* and Lada A. Adamic

*Correspondence:

shankar94@fb.com

Core Data Science, Facebook, 1
Hacker Way, Menlo Park,
California, USA

Full list of author information is
available at the end of the article

Abstract

In the classic “influence-maximization” (IM) problem, people influence one another to adopt a product and the goal is to identify people to “seed” with the product so as to maximize long-term adoption. Many influence-maximization models suggest that, if the number of people who can be seeded is unconstrained, then it is optimal to seed everyone at the start of the IM process. In a recent paper, we argued that this is not necessarily the case for social products that people use to communicate with their friends [1]. Through simulations of a model in which people repeatedly use such a product and update their rate of subsequent usage depending upon their satisfaction, we showed that overambitious seeding can result in people adopting in suboptimal contexts, having bad experiences, and abandoning the product before more favorable contexts for adoption arise. Here, we extend that earlier work by showing that the costs of overambitious seeding also appear in more traditional threshold models of collective behavior, once the possibility of permanent abandonment of the product is introduced ~~in addition to the model of repeated-product usage that we studied in reference [1]~~. We also further demonstrate that these costs can be mitigated by using conservative seeding approaches besides those that we explored in the earlier paper. Synthesizing these results with other recent work in this area, we identify general principles for when overambitious seeding can be of concern in the deployment of social products.

Keywords: social networks; influence maximization; diffusion on networks

1 Introduction

The study of how new ideas and products spread through networks dates back decades, to early studies from the 1950s and 1960s of the adoption of health innovations [2, 3, 4] and to the development of general models of product adoption by Rogers and Bass [5, 6]. An important milestone was the formulation of “influence maximization” (IM) as an algorithmic problem by Domingos and Richardson [7]. In IM, a product developer typically has limited resources (e.g., an advertising budget) with which to give or market a product to potential adopters. The developer assumes that adoption of the product spreads through the social network of potential adopters through some peer-influence process. Then, the challenge is to decide which people to “seed” with the product in order to maximize long-term adoption. Since its formulation by Domingos and Richardson, the influence maximization problem has found applications across diverse domains, from traditional applications in marketing [8], to the spreading of health information [9, 10], to the diffusion of microfinance programs in villages [11].

Influence maximization has been theoretically studied under a variety of peer-influence models. One classic IM model is the independent-cascade model, in which

friends of new adopters also adopt with some probability [12]. So-called linear threshold models comprise another class, in which people will adopt if sufficiently many of their friends adopt [13, 14]. Depending upon the specific threshold model, an individual’s adoption decision can depend upon a minimum number of friends adopting or upon a minimum percentage of friends adopting [15]. Soon after the formulation of IM by Domingos and Richardson, Kempe, Kleinberg, and Tardos demonstrated both that IM is NP-hard under the usual independent cascade and threshold models and that there are nevertheless simple greedy algorithms for selecting the seeds with strong performance guarantees [16]. Their work has inspired a large literature around developing even better heuristic algorithms for IM. A recent review of state-of-the-art algorithms can be found in reference [17].

In this paper, we revisit a question that we previously explored [1]: if there is no budgetary constraint on seeding, is it optimal to seed everyone at the start of the IM process? Despite the general hardness of IM, the traditional independent cascade and threshold models all agree that the answer to this question is “yes.” Does that property of these simplified models provide reasonable guidance for real product-deployment scenarios? There are several reasons why it may not, including costs associated with people rejecting the product, downstream word-of-mouth effects, and so-called “non-conformism” effects, where people are inclined to adopt *less* popular products. We review prior work on each of these pathways to “overexposure” in Section 2 below.

Our main focus here, however, is on a distinct pathway to overexposure, which we demonstrated in a recent paper [1]: when the product under consideration is one that allows people to communicate with their friends (i.e., a “social” product), if people adopt too early, then they may begin using the product in contexts where insufficiently many of their friends are using it. This can lead to abandonment of the product prior to the emergence of a more favorable context for adoption. In our earlier paper, we showed that a more conservative seeding strategy can often help avoid these premature abandonments of the product and lead to greater long-term usage. Crucially, we showed that this remains true even in the absence of a budgetary constraint on initial seeding: even if a product developer *can* simply hand the product to *everybody*, it may be preferable not to do so.

The present paper extends our previous work in various ways. In reference [1], we demonstrated the “costs” of overambitious seeding in a model of repeated product usage, where people gain access to a social product and then either use the product or abstain in a sequence of time steps. Here, we show that these “costs” also appear in more traditional threshold models, once the possibility of permanent abandonment of the product (also referred to as “churn”) is introduced. Furthermore, through simulations on networks with a clear community structure, our earlier work showed that seeding approaches that focus on one of the clusters can often outperform approaches that seed the entire network. In this paper, we show that there are conservative seeding approaches that do not rely on clear-cut community structure, but which still lead to greater longer-term adoption than universal seeding. After demonstrating the robustness of our previous results in these two different ways, we then attempt to abstract away general principles for when product developers ought to factor these considerations into their product deployment decisions.

The rest of this paper is structured as follows. Section 2 places our work in the context of prior research on overexposure in IM. With this context in place, in Section 3, we introduce the repeated-usage and threshold models that we study in this paper. Next, in Section 4, we study each of these models on certain, very special network structures, developing intuitions for *why* overambitious seeding can be problematic in both models. In practice, of course, we will want to see how the models behave on more realistic network structures, and to that end, in Section 5 we introduce the real-world network structures that we use in our numerical simulations. Sections 6 and 7 then report our simulation results, showing how two different conservative seeding approaches can outperform universal seeding. In Section 8, we extract some general principles for when overambitious seeding can be costly before concluding in Section 9 by reviewing our findings and pointing out opportunities for extensions.

2 Related Work

In this section, we review research on overexposure and overambitious seeding in influence maximization. Our goal here will be to examine the implications of various previously explored models for the fundamental question articulated above: in the absence of a budgetary constraint on the seeding process, is it optimal to seed everyone immediately? This survey of prior research helps distill the reasons why it is interesting that, in each of the models studied in this paper, the answer to this question is often “no.”

In [16], Kempe et al. showed that the classic independent cascade and linear threshold models obey a monotonicity property, where a subset of a cohort of initial adopters cannot lead to higher long-term adoption than the entire cohort. Furthermore, they generalized these models to a larger class of so-called “triggering” models, in which each subset of a person i ’s neighbors is associated with a probability of i adopting, and showed that triggering models also exhibit monotonicity [16]. If this monotonicity property holds, and if each person accepts the seed independently, then the optimal approach in the unbudgeted case clearly involves seeding everybody: an unseeded individual’s probability of adoption in the seed round is 0, and by monotonicity, it would be preferable if that probability was non-zero.

Models of overexposure generally try to show that there are plausible assumptions about real-world IM settings that can violate monotonicity. One path to overexposure involves introducing some type of negative payout for rejection of the product. A recent example of this is reference [18], in which Abebe, Adamic, and Kleinberg study a diffusion process where there are positive payouts for adopters and negative payouts for rejecters. If someone adopts, that person will refer the product to his or her friends, which could result in further adoptions or rejections. In this model, there can be circumstances where it is detrimental to seed an individual i , because the costs of the product being exposed to i ’s friends may outweigh the benefits of i ’s adoption. In the unbudgeted case, the optimal strategy still does not generally involve seeding everyone, because that would expose the product to many rejecters, leading to potentially avoidable negative payouts [18]. The results of Abebe et al. echo empirical findings such as the so-called Groupon effect, where exposure to a larger audience can have unintended negative effects (e.g., upon Yelp ratings) [19].

Other authors have studied overexposure effects arising from more direct negative externalities of adoption, such as “negative word-of-mouth” [20]. Empirical research actually suggests that dissatisfied adopters spread their perspective more often than satisfied adopters, sharing their negative sentiment with up to ten friends [21]. Cui et al. recently reported results for a model where satisfied adopters can enhance the probability of subsequent adoption by their friends, while dissatisfied adopters can reduce that probability [22]. In such a setting, it may be preferable to seed people who are likely to spread positive word-of-mouth and avoid seeding others.

“Non-conformist” or “hipster” effects comprise yet another class of negative externalities. “Hipsters” in these models refrain from adopting products that are too popular and / or abandon products if they become too popular [23]. Although not strictly framed as an IM study, the recent work of Juul and Porter shows how the presence of hipsters can have dramatic effects upon the long-term adoption of two competing products. Indeed, in some cases, the product that begins the process with no adopters at all ends up accounting for the majority in the steady state [24].

Kempe, Kleinberg, and Tardos referred to models in which adopters can revert to the non-adopting state as “non-progressive” models, to contrast with “progressive” models where people can only transition *into* the adopting state. If we want to consider the product experiences of people after they make their initial adoption decision, then some form of non-progressive model is appealing. Kempe et al. showed that the simplest non-progressive extensions of monotonic triggering models (e.g., where people abandon the product if enough of their friends do) inherit the monotonicity property. This is because these models can be mapped to their progressive counterparts on a temporal network in which people are represented by a node in each temporal layer, and there are links between each person i at time t and their friends at time $t - 1$ ^[1] [16]. This argument for monotonicity does not work if the original progressive model is itself non-monotonic, or if people abandon the product permanently after a fixed number of adoptions.

In reference [1], we previously argued that it can be detrimental to seed everyone at the start of an unbudgeted IM process in a certain type of non-progressive model, even in the absence of the mechanisms studied in the previous literature surveyed above. Our model was motivated by “social products” that are used by friends to communicate with one another, and it considered the product experiences of people after adoption instead of focusing exclusively upon the binary adoption / rejection process. A key point of our earlier paper was that taking into account these product experiences naturally leads to the emergence of costs of overambitious seeding, even in the absence of negative payouts of rejection, negative word-of-mouth, and non-conformism effects. However, reference [1] made this point in a model that is rather structurally different from classic models of IM. Here, we show that the

^[1]If such a mapping holds and the non-progressive model is monotonic, it still may make sense to employ a gradual seeding approach in a budgeted scenario. Indeed, Jankowski et al. have recently explored the benefits of gradually seeding parts of the network that have not been activated by previous seeding rounds, instead of seeding all at once and possibly wasting resources on parts of the network that would have adopted anyway [25]. Note, however, that the notion of “wasting” seeding resources in that work depends upon the existence of a budget.

same mechanisms can lead to costs of overambitious seeding in traditional threshold models, once the possibility of permanent churn is included.

3 Models of Social-Product Usage

In this section, we introduce two models of how people adopt, use, and abandon social products. First, we review the model of repeated product usage that was proposed in reference [1], which considers the gradual impact of individual product experiences upon people’s subsequent behavior. Then, we propose a modification of the traditional threshold model as a “coarser-grained” model of long-term adoption decisions.

Both models are intended to describe the choices of people embedded in an undirected social network. Each node i represents a person who can potentially use or adopt the product. Each edge ij represents a friendship tie between two people.

Repeated-Usage Model: Our repeated-usage model proceeds in a sequence of time steps, beginning with $t = 0$. At any time t , a person i can either have access to the product or not. If the person has access, then the person uses the product in that time step with probability $p_i(t)$ and abstains otherwise. At the time t_i when a person i initially gets access, $p_i(t_i)$ is initialized to a value p_0 .

We associate a threshold s_i with each person. If i uses the product in time step t , then s_i is the number of friends of i who also need to use the product at time t for i to be satisfied. Then, i adjusts his or her probability of subsequent usage up or down as follows:

$$p_i(t+1) = \begin{cases} p_i(t) + \delta & \text{if } \color{red}{> s_i} \text{ more than } s_i \text{ friends use in time step } t \\ p_i(t) & \text{if exactly } s_i \text{ friends use in time step } t \\ p_i(t) - \delta & \text{if } \color{red}{< s_i} \text{ fewer than } s_i \text{ friends use in time step } t \end{cases} \quad (1)$$

We allow $p_i(t)$ to grow to 1 or drop to 0. While $p_i(t) = 1$ is not necessarily a permanent state, $p_i(t) = 0$ is permanent, because ~~the person will no longer have any product experiences. Consequently, the person will no longer have any~~ it guarantees that the person will no longer have any product experiences, and consequently, will have no opportunities to increment their usage.

In this model, in situations where we do not give access to everyone at time $t = 0$, we need some protocol for implementing the gradual expansion of access. As in [1], we expand access to a new person when they have had at least two friends using the product in each of five consecutive time steps. This is one example of a more conservative seeding strategy than universal seeding at $t = 0$. Other variants of this rule can certainly be considered and may even lead to better long-term outcomes, but this choice suffices to demonstrate our main results.

Threshold Model with Churn: Various types of threshold models have been proposed in a number of contexts, from models of percolation in statistical physics [26] to collective models of social behavior [13, 15]. When these models are used to study the decisions of people situated in a social network, the general idea is that people are able to take on one of two states, which we can refer to as “adopting” and “non-adopting.” Certain people begin in the adopting state (e.g., through the outcome of a seeding process). Then, others may adopt if sufficiently many of their

friends are in the adopting state. The adoption rule may be formulated in terms of the absolute number of friends, or alternatively, it may be formulated in terms of a percentage of friends. Generally, the adoption rule is iteratively applied until no more people would adopt.

In “non-progressive” threshold models, people can also transition out of the adopting state. For example, as an outcome of the seeding process, some people can find themselves in a situation where insufficiently many of their friends are adopting. In these circumstances, they may transition back to the non-adopting state. This can, in turn, leave others in a situation where they have too few adopting friends, leading to more defections. These transitions out of the non-adopting state will, in general, cooccur and compete with transitions *into* the adopting state over time. If people are willing to adopt the product an arbitrary number of times, then the non-progressive model can be mapped to a progressive model and is monotonic in the size of the original seed set [16]. However, if people permanently churn after a fixed number of adoptions, then the non-progressive model is not necessarily monotonic.

The threshold model that we study here proceeds as follows:

- **Seeding Stage:** At time $t = 0$, certain people within a social network are offered the product, which they adopt with acceptance probability p_a .
- **State Updates:** At each subsequent time step $t = 1, 2, 3, \dots$, people update their states in two successive waves, which continue until the process converges:
 - 1 **Adoption Round:** People who are not currently adopting look at the states of their friends after the ~~last~~previous churn round (at time step $t - 1$) and adopt if at least s_i of their friends are adopting.
 - 2 **Churn Round:** People who are currently adopting look at the states of their friends after the ~~last~~previous adoption round (at time step t) and churn if fewer than s_i of their friends are adopting.
- **Constraints on State Changes:** The state updates described above are constrained by the following two rules:
 - 1 **One-Time-Step Commitment:** People who adopt in an adoption round do not churn in the immediately following churn round. There is a rate limit to these state changes because we are modeling long-term changes in people’s attitudes towards the product.
 - 2 **Single Adoption per Person:** People give the product only one chance before churning permanently.

Comparing the Two Models: In [1], we motivated the repeated-usage model through the following assumptions about social product usage:

- 1 **Need for social support:** a person’s satisfaction with a product experience depends upon how many of their friends are using it.
- 2 **Rate-of-usage adjustments:** When people gain access to the product, they begin using it at a low rate p_0 and then gradually ramp their rate of usage up or down depending upon whether they are satisfied with their experiences.
- 3 **Possibility of permanent churn:** if people have enough unsatisfying product experiences, they churn permanently and are unwilling to try the product again.

Our threshold model also clearly satisfies the “need for social support” assumption and, like the repeated-usage model, encodes this property through the parameters s_i . Moreover, the threshold model satisfies the “possibility of permanent churn” assumption. The threshold model does not, however, incorporate gradual “rate-of-usage adjustments” but, rather, binary state changes between adoption and non-adoption. This is in keeping with its being a temporally coarse-grained model of adoption decisions.

In [1], we also emphasized that the repeated-usage model *excludes*:

- 1 a budgetary constraint on seeding
- 2 rejection of the seed
- 3 negative word-of-mouth or nonconformism effects

Our threshold model also excludes budgetary constraints, negative word-of-mouth, and nonconformism effects. However, when $p_a < 1$, we do allow rejection of the seed. Since there are no negative externalities to adoption in our model, there can be no costs to overambitious seeding if the seeding process is universally successful. We will show, however, that costs naturally emerge if the success of the seeding process is stochastic. Still, there is no direct cost to someone rejecting the seed, so the path to overambitious seeding here is distinct from the one explored, for example, in [18].

Comparing the roles of the parameters p_0 and p_a in the repeated-usage and threshold models respectively, we can refine our fundamental question for each context. In the repeated-usage model, by fixing a low p_0 , we pose the question: is it optimal to seed everyone at time $t = 0$ if every seeded person adopts, but subsequently uses at a low rate? Meanwhile, by fixing a low p_a in the threshold model, we pose the question: is it optimal to seed everyone at time $t = 0$ if seeding succeeds only at a low rate? The simulation results of Sections 6 and 7 show that the answer to both of these questions is often “no.”

4 Toy Examples

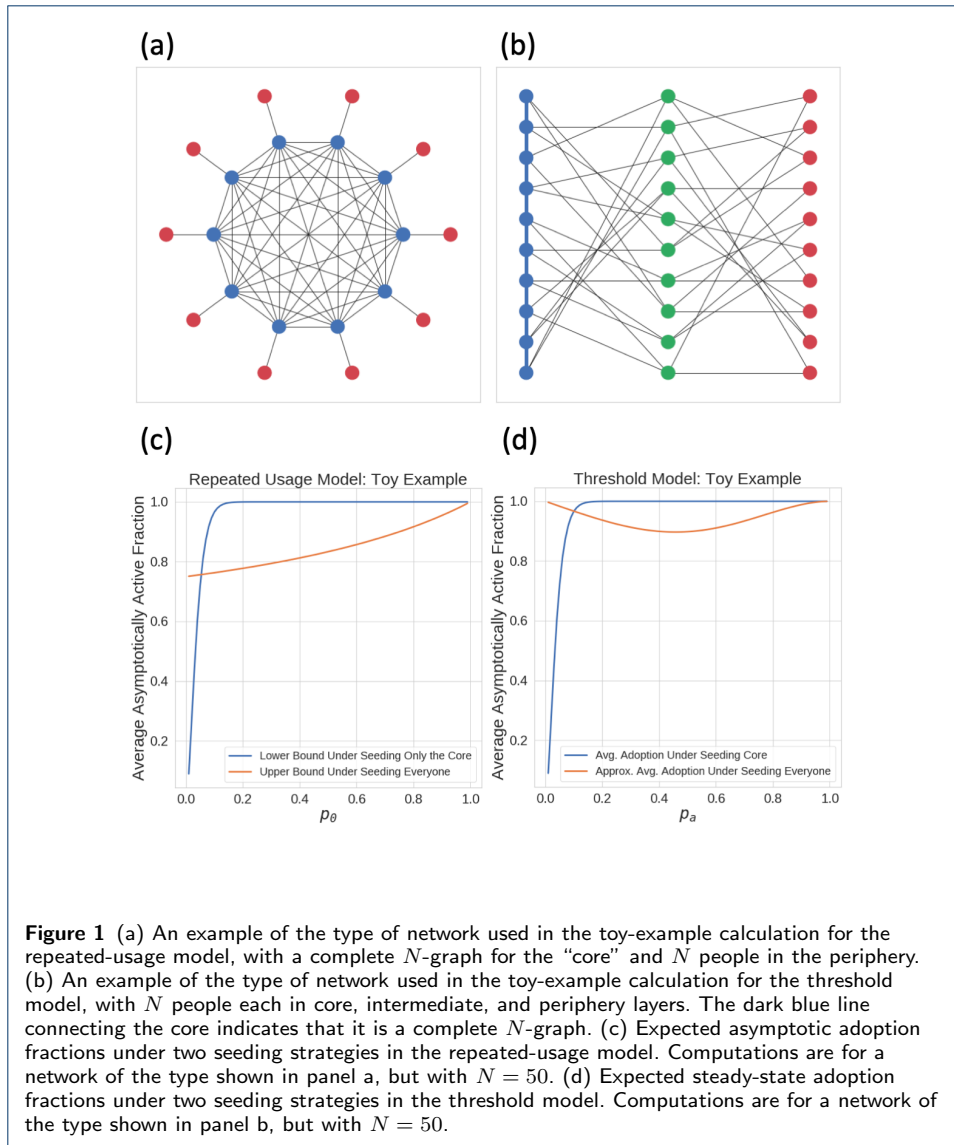
Before proceeding to the simulation results, we dedicate this section to analytical investigation of our models on certain, very special network structures. These “toy” examples illustrate the mechanisms through which overambitious seeding can reduce long-term adoption. Then, the simulation results of subsequent sections show that these mechanisms are relevant in more general contexts.

Repeated-Usage Model: Suppose we run the repeated-usage model on a network with a very strong core-periphery structure [27]. In particular, consider a situation where the “core” consists of a complete N -graph (i.e., N people who are all friends with the $N - 1$ others) and the “periphery” consists of N people, each of whom is friends with one person in the core. An example of such a network is shown in Figure 1a.

As usual in the repeated-usage model, rates of usage $p_i(t)$ are initialized to p_0 when people receive access. However, for the present purposes, suppose that update rule for $p_i(t)$ is a much simpler variant of the one proposed in equation (1):

$$p_i(t+1) = \begin{cases} 0 & \text{if no friends use in time step } t \\ 1 & \text{if at least 1 friend uses in time step } t \end{cases} \quad (2)$$

The simplified update rule (2) has certain important corollaries:



- 1 **Satisfaction is reciprocal:** because everyone only needs one active friend to be satisfied, if a person i is satisfied with a product experience, then so are all of i 's active friends.
- 2 **The state where $p_i(t) = 1$ is permanent:** this follows from item 1, because any i with $p_i(t) = 1$ has at least one friend j with $p_j(t) = 1$. This implies that all of i 's subsequent product experiences will be satisfying.
- 3 **Churn is only possible on the first product experience:** this follows from item 2, because if a person i does not churn on the first product experience, then that person ends up in the state where $p_i(t) = 1$.

When coupled with the special network structure of Figure 1a, there is another important implication: once one person in the core has had a satisfying experience, then everyone who is subsequently active in the core will be satisfied.

First, we consider the case where we only give access to the core at time $t = 0$. The probability that at least two people are active in the core in the first time step

is:

$$1 - (1 - p_0)^N - Np_0(1 - p_0)^{N-1} \quad (3)$$

If this occurs, ~~then all of the active people then those two people~~ will be satisfied with their experience and update their rates of usage to 1. Then, in the ensuing time steps, others in the core will try out the product, have satisfying experiences, and update their rates of usage to 1 as well. ~~This process will take some time, since each person's initial choice to be active is independent and will take time $\frac{1}{p_0}$ on average. However,~~ if we wait until everyone in the core is consistently active, we can then grant access to the periphery in circumstances where all people in the periphery are guaranteed to have satisfying experiences. Thus, equation (3) is a lower bound on the probability with which we can end up with all $2N$ people adopting. This actually gives a very conservative underestimate of the average adoption fraction^[2], but the bound implied by equation (3) is sufficient to demonstrate our main point.

To see why, let us now consider the case where we grant access to everyone at time $t = 0$. Here, we can lower bound the probability that a person i in the periphery will churn. In particular, we can bound it by the sum over all times T of the probability that i is first active in time step T and that i 's friend in the core is not active at all up to and including time step T ^[3]. This gives:

$$\sum_{T=0}^{\infty} (1 - p_0)^{2T+1} p_0 = \frac{p_0(1 - p_0)}{1 - (1 - p_0)^2} \quad (4)$$

Then, we can upper bound the average long-term adoption fraction by assuming that everyone except this fraction adopts:

$$1 - \frac{p_0(1 - p_0)}{2(1 - (1 - p_0)^2)} \quad (5)$$

We compare equation (3) to equation (5) in Figure 1c. This shows that a lower bound on the adoption under seeding the core beats an upper bound on the adoption under seeding everyone over a broad range of values of p_0 . The reasoning above exposes why this is the case: by granting access to the periphery too early, we expose people in the periphery to the product before they are likely to be satisfied with their product experiences. Furthermore, except at very low p_0 , this premature

^[2]Because we have not added in possible scenarios where, for instance, no one in the core is active in the first time step but two are active in the second, where one person in the core churns in each of the first two time steps but two people are active in the third, etc. In each of these cases, a very large fraction of the population can nevertheless be active asymptotically.

^[3]Here, we are neglecting situations where i 's friend in the core already churned due to having an unsatisfying experience before time T . This is a small effect as N gets large because of the low likelihood of having bad experiences in the core. Meanwhile, if p_0 is small, it takes each person $\frac{1}{p_0}$ time steps on average to be active at all, so it is quite likely that a specific individual in the core is inactive at early time T . This is the effect that we capture in equation (4).

exposure of the periphery confers no benefit to the core, which is sufficiently dense to produce satisfying experiences all on its own. It is better to wait until the core is activated and only then to grant access to the periphery.

Threshold Model with Churn: Next, we consider running our threshold model on the network shown in Figure 1b [27]. This is a network with a dense “core” of N people who are all connected to one another and who are represented by the blue nodes. There is an “intermediate” layer of N people who are each friends with two randomly chosen people in the core; the people in this “intermediate” layer are represented by the green nodes. Finally, there is a “periphery” of N people who are friends with two randomly chosen people in the intermediate layer and who are represented by red nodes. We will consider a case where $s_i = 2$ for all people in the network and where p_a can vary. We will then ask whether it is better, in terms of asymptotic adoption, to seed everyone or to only seed the core.

First, we consider the case where we only seed the core. The probability that fewer than 2 people in the core adopt under seeding is given by:

$$(1 - p_a)^N + Np_a(1 - p_a)^{N-1} \quad (6)$$

With this probability, adoption dies out completely. On the other hand, if at least 2 people adopt under seeding, then by time $t = 1$, the entire core will adopt. Potentially, some people in the intermediate layer will as well, if they happen to have two friends in the core who adopted during the seeding round. By time $t = 2$, the entire intermediate layer will adopt, because every person in the intermediate layer has two adopting friends in the universally adopting core. Potentially, some people in the periphery will adopt as well, if they happen to have two intermediate layer friends who adopted by time $t = 1$. Finally, by time $t = 3$, the entire periphery will adopt as well, because every person in the periphery has two adopting friends in the universally adopting intermediate layer. Thus, as long as 2 people in the core adopt during the seeding round, the entire network eventually adopts. This means that the average final adoption fraction is:

$$1 - (1 - p_a)^N - Np_a(1 - p_a)^{N-1} \quad (7)$$

Next, we turn to the case where we seed everyone. The probability that at least two people in the core adopt under seeding remains the same. If that happens, then no one in the intermediate layer will churn, including those who happened to adopt under seeding. This is because, as of the adoptions that occur at $t = 1$, every person in the intermediate layer will have 2 adopting friends in the core. However, people in the periphery who adopt under seeding *can* churn. For a $t = 0$ adopter in the periphery to *not* churn, one of the following must be true by the first churn round:

- 1 Their two friends in the intermediate layer adopted under seeding.
- 2 One of their friends in the intermediate layer adopted under seeding, and the other had two friends in the core who adopted under seeding.
- 3 Neither of their friends in the intermediate layer adopted under seeding, but both had two friends in the core who adopted under seeding.

In principle, in case 3, the four friends-of-friends in the core need not all be distinct; however, as N grows very large, we can ignore this possibility. Then, the probability that a person in the periphery churns is approximately:

$$p_a [1 - p_a^2 - 2(1 - p_a)p_a^3 - (1 - p_a)^2 p_a^4] \quad (8)$$

Hence, in the case where at least two people in the core adopt upon seeding, we can expect the final fraction of adopters to approximately be:

$$\frac{2}{3} + \frac{1}{3} [1 - p_a + p_a^3 + 2(1 - p_a)p_a^4 + (1 - p_a)^2 p_a^5] \quad (9)$$

When fewer than two people in the core adopt upon seeding everyone, it is still possible for long-term adoption to be sustained if sufficiently many people adopt in the intermediate and periphery layers. This exemplifies how seeding everyone can sometimes be beneficial, especially at small values of p_a . Nevertheless, when $p_a \gg \frac{2}{N}$, equation (9) will still be a good approximation to the average adoption [4]. We compare equations (7) and (9) in Figure 1d. Here too, we see that it is preferable to seed only the core over a large range of p_a . Yet again, the costs of overambitious seeding originate in the premature exposure of the people in the periphery to the product. These people's abandonment of the product is avoidable under a more conservative seeding strategy that focuses on the densest part of the network.

Comparing Figure 1d to Figure 1c, we see that the costs of overambitious seeding are maximized at some intermediate value of p_a for the threshold model, while these costs get bigger as p_0 gets lower in the repeated-usage model. This is due to p_a playing a dual role in the threshold model, determining both the proportion of the population that is exposed early and the average social support that population can expect. We will return to this point in Section 8, when we discuss general settings in which overambitious seeding can be especially problematic.

5 Networks Used in Simulations

In [1], to argue that overambitious seeding can be problematic on real-world networks, we ran simulations of the repeated-usage model on portions of the Facebook friendship graph, known as Social Hash (SH) clusters. The SH clustering was originally developed to enable faster data retrieval by physically collocating data for people who communicate frequently. Thus, many (but not all) of a person's frequently contacted friends belong to the same SH cluster [28, 29].

This property is well matched to the type of cluster-level approaches that we tested previously [1], and the same is true here. Therefore, in this paper as well,

^[4]Note that equation (9) is a poor approximation to the adoption fraction when $p_a \approx \frac{2}{N}$ or lower for at least three reasons. First, we need to incorporate corrections to (8) arising from the fact that we condition on at least two people in the core adopting under seeding in that calculation. Second, to produce a good estimate in this regime, we cannot neglect the case of fewer than two core adopters accurately. Third, a good approximation in this regime must approach zero adoption as p_a goes to zero; equation (9) does not exhibit this behavior.

we will report simulation results modeled on de-identified SH clusters containing US Facebook users who visited in a 28 day period. When we discuss cluster-based seeding of the repeated-usage model (in Section 6), we report data from simulations on SH clusters computed on 2018-04-29. The properties of the SH clusters used in these simulations can be found in [1]. All other simulations were performed on SH clusters computed for active US users who visited in the 28 days leading up to 2019-01-27.

As in reference [1], we also select three-cluster networks such that each cluster has average out-of-cluster degree $\langle k_{oc} \rangle \geq 1$ ^[5]. Tables 1 and 2 report statistics of the distributions of the within-cluster degree k_{ic} and out-of-cluster degree k_{oc} for the various SH clusters and three-cluster networks. These tables show that there is considerable structural diversity amongst these clusters and networks^[6].

Figure 2 shows an example of a three-cluster SH network. This is the network qrs from Table 2.

6 Simulation Results: Cluster-Based Seeding

Previously [1], we demonstrated the costs of overambitious seeding in the repeated-usage model by showing that seeding a single cluster can lead to greater long-term adoption than seeding all three clusters in a variety of SH networks. In this section, we recap the results of [1] for the repeated-usage model and then move on to show that the same phenomenon can be observed in the threshold model as well, albeit in a quantitatively weaker form.

Repeated-Usage Model: Figure 3 shows simulation results for the repeated-usage model on three-cluster SH networks. These are the three-cluster networks that we introduced in reference [1], and they can be distinguished from the newer clusters used in subsequent sections by the use of uppercase letter labels. In these simulations, we fix $s_i = 2$ for all people, vary p_0 , and ask which of the following strategies leads to most long-term adoption:

- 1 Seed the cluster with the highest median within-cluster degrees k_{ic} .
- 2 Seed the two clusters with the highest median within-cluster degrees k_{ic} .
- 3 Seed all three clusters.

^[5]On average, each person in each cluster has at least one out-of-cluster friend.

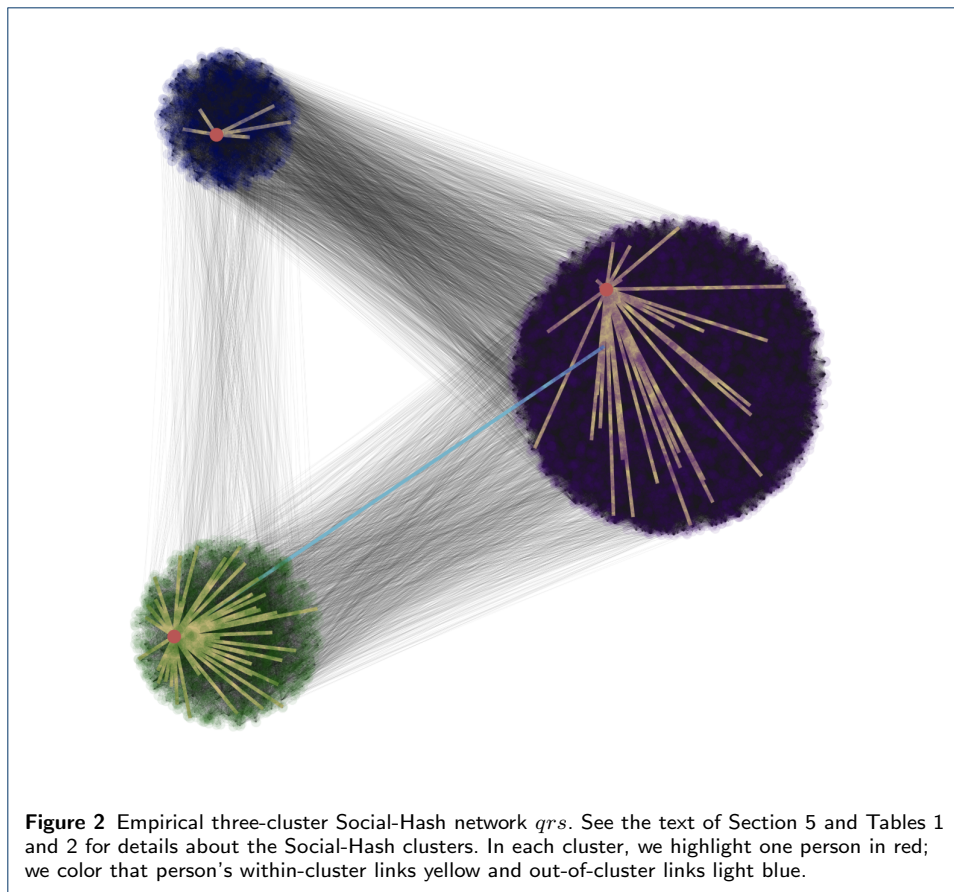
^[6]Table 1 reports structural properties for 25 SH clusters. ~~The first 10 of these~~ Clusters $a-j$ were sampled for the purpose of comparing k -core seeding and universal seeding in the two models. We did not end up reporting results for clusters c and g in Figures 5 and 6 because these clusters do not have a 10-core, so they are excluded from Table 1. In Figure 5, we also did not run simulations for clusters b , d , and j because the repeated-usage model is expensive to simulate, and these clusters ended up being too large. ~~The last 15 clusters in Table 1~~ Clusters $k-y$ were sampled for the purpose of comparing cluster-based seeding approaches in the threshold-model. These clusters form parts of three-cluster networks whose properties are reported in Table 2. The size of the three-cluster network can differ slightly from the size of the three clusters individually because, in both cases, we exclude people with zero degree, who would inevitably churn under our model. In a small percentage of cases, a person who has no within-cluster friends may still have friends in another cluster when three clusters are considered together.

SH Cluster Properties								
ID	N_c	$\langle k_{ic} \rangle$	P5 k_{ic}	P50 k_{ic}	P95 k_{ic}	$p_{10-core}$	C	r
a	1214	12.8	1	8	39	28.75%	0.20	0.17
b	52985	123.6	3	58	464	84.62%	0.16	0.07
d	40643	119.0	1	56	454	79.64%	0.27	0.16
e	1868	14.0	1	9	42	29.23%	0.24	0.37
f	13726	51.6	4	40	135	85.56%	0.20	0.24
h	7306	11.5	1	8	34	18.30%	0.30	0.16
i	4919	29.9	2	18	100	64.73%	0.18	0.12
j	36238	94.8	1	44	362	78.27%	0.33	0.15
k	4797	124.9	4	75	429	88.60%	0.23	0.12
l	2717	12.4	1	6	43	28.34%	0.07	0.16
m	3546	20.7	2	14	62	55.89%	0.11	0.24
n	6605	8.0	1	5	25	4.19%	0.22	0.27
o	4223	8.0	1	5	25	3.29%	0.22	0.10
p	1139	9.8	1	5	34	22.04%	0.23	0.10
q	4473	105.6	1	58	349.4	68.99%	0.45	0.35
r	851	37.2	1	4	151.5	40.54%	0.40	0.23
s	1561	18.0	1	8	66	38.24%	0.26	0.34
t	13194	33.8	4	21	111	75.62%	0.27	0.28
u	7576	37.7	3	20	133	70.67%	0.20	0.16
v	3562	18.9	1	11	61	48.74%	0.31	0.07
w	986	8.9	1	6	25	3.45%	0.23	0.18
x	3085	43.6	2	28	137	76.11%	0.24	0.29
y	11349	36.3	1	21	126	66.27%	0.28	0.22

Table 1 Statistics of individual SH clusters. Here, k_{ic} refers to the within-cluster degree, $p_{10-core}$ is the percentage of people who belong to a 10-core when the cluster is considered in isolation, C is the clustering coefficient, and r is the degree assortativity. **In response to a comment from reviewer 1, we have dropped unused clusters 3 and 4 from this table**

SH Network Properties						
ID	N	Percentage of people whose k_{oc}				
		= 0	= 1	= 2	= 3	> 3
klm	11158	19.6%	12.4%	9.0%	6.7%	52.3%
nop	12224	45.5%	22.4%	11.9%	7.1%	13.1%
qrs	6891	68.3%	10.8%	5.5%	3.1%	12.2%
tuv	24406	36.9%	18.6%	10.9%	7.7%	25.9%
wxy	15447	57.6%	14.8%	8.0%	4.6%	15.0%

Table 2 Statistics of networks composed of three SH clusters from Table 1. Here, k_{oc} is the out-of-cluster degree.



We report the average adoption in the last 100 time steps of 10000 time step simulations. Simulation results from reference [1] showed that 10000 time steps are generally sufficient for adoption to reach its asymptotic value.

The left-hand panels of Figure 3 show the fraction of the network that has access at late times, and the right-hand panels show the fraction that is active. In these five three-cluster networks, we see that seeding one cluster consistently beats seeding three clusters in the low p_0 regime. The “costs” of overambitious seeding here can be substantial: for networks LMN , OPQ , and RST , we observe up to 35-45% less activity under the universal seeding strategy. There is only one network (FGH) where universal seeding at $t = 0$ ever wins, and then only at high p_0 . Meanwhile, the left-hand panels of Figure 3 show that, when the single-cluster seeding policies win, it is often despite the fact that there are people who are *never* given access.

We put forward an argument for *why* seeding a single-cluster is so often preferable in reference [1], which goes as follows: at early times in the repeated-usage model, we are faced with a fundamental tradeoff. There are costs to seeding a cluster, because by assumption, the rate of initial activity p_0 is low. Therefore, some people will adopt in unfavorable contexts, meaning that they will typically be unsatisfied by how many of their friends are active when they are. This will result in some permanent churn. On the other hand, there are also costs to *not* seeding a cluster: in particular, people in other clusters lose out on the social support of people in the unseeded cluster. We saw the same tradeoff at play in the toy example of Section 4.

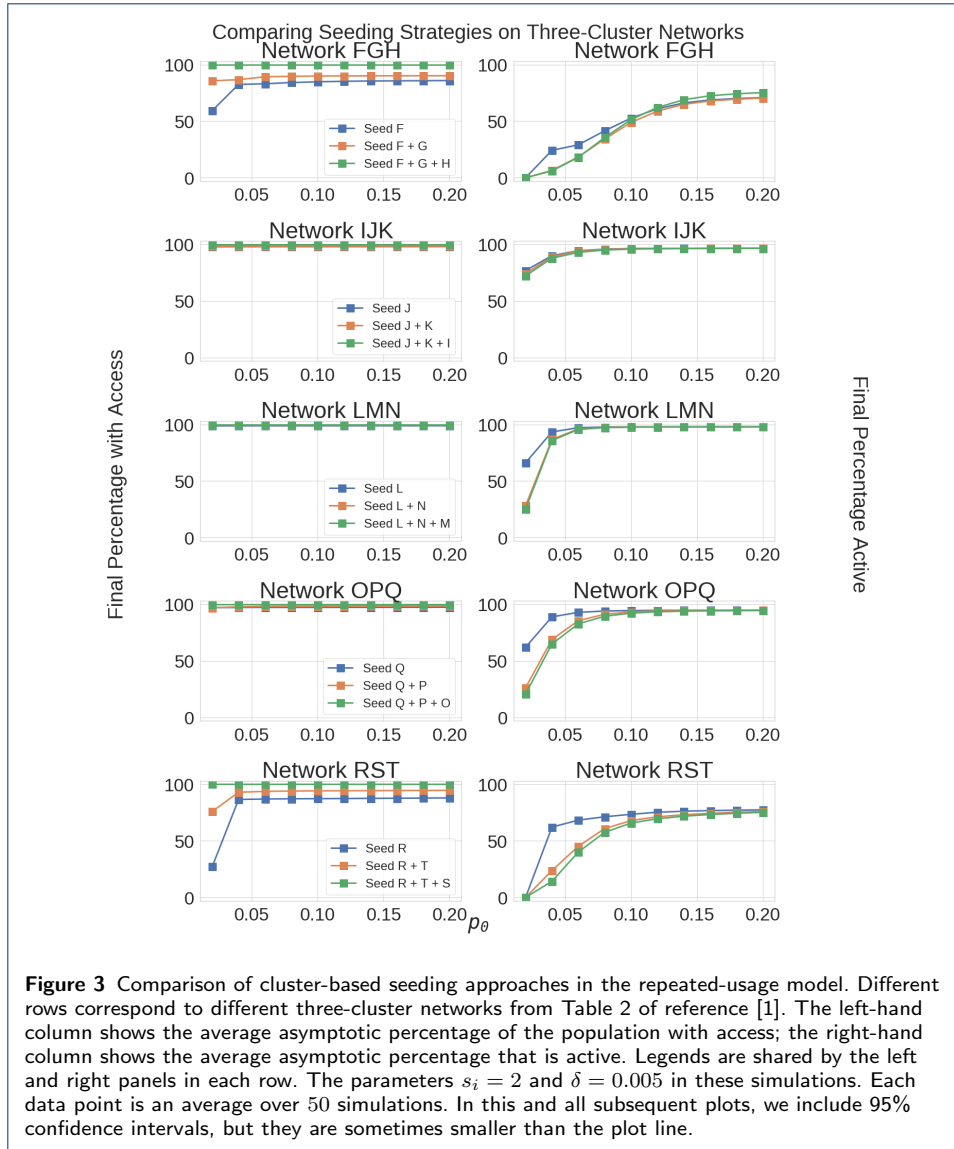
In reference [1], through simulations on synthetic networks, we demonstrated multiple regimes where this tradeoff plays out in different ways. When p_0 is very low, activity is not sustained in the long-term under any seeding strategy. As p_0 is tuned up from this regime, we initially enter a regime where the combined early activity in all three clusters is sufficient to sustain long-term activity (i.e., the universal seeding policy wins). At higher values of p_0 , two clusters can sustain long-term activity in isolation, and it is “costly” in terms of asymptotic activity to seed the third. In other words, seeding the third cluster at time $t = 0$ results in churn that could have been avoided by waiting and granting access to the third cluster under more favorable circumstances, when people in the two seeded clusters are active at very high rates. Finally, if p_0 is sufficiently high, a single cluster can sustain long-term activity in isolation, and it is costly to seed any more at $t = 0$.

When studying this model on empirical networks, we typically *only* observe the final regime, because of the inherent heterogeneity in the within-cluster degree distribution. If we seed a single cluster, there is usually some subnetwork of that cluster (e.g., perhaps involving the highest-degree people) where long-term activity can build up in isolation. Then, the activity in that subnetwork is usually sufficient to bootstrap the spreading of favorable contexts for adoption through the rest of the three-cluster network. Note that the other two regimes (where it is preferable to seed 2 or 3 clusters) presumably still exist; we just do not observe them in Figure 3 because they occur in a very narrow and low range of p_0 . Furthermore, idiosyncrasies of graph topologies in empirical networks can produce cases like *FGH*, where at high p_0 we reenter a regime where seeding all three clusters is preferable. Despite these anomalies, seeding one cluster very generally beats seeding three when the repeated-usage model is simulated on real-world networks.

Threshold Model with Churn: In the case of the threshold model, simulations are efficient enough that we can simulate every possible cluster-based seeding strategy for each of the five three-cluster networks. We simulate a case where $s_i = 5$ for all people. The final adoption curves vs. seed acceptance probability p_a are shown in Figure 4 .

The clearest case here is network *nop*. At the lowest values of p_a , the strategy of seeding all clusters leads to the highest asymptotic activity. This is for the same reasons that we discussed in the case of the repeated-usage model: there is a tradeoff at early times between exposing people to the product prematurely and missing out on the social support that these people could provide to others. Asymptotic activity first develops when the combined early activity in all three clusters is sufficient to sustain long-term usage. As p_a increases though, we observe a small regime (around $p_a = 0.08$) where seeding two clusters is optimal. Finally, we enter a regime where seeding just one cluster (cluster *p*) beats out all other strategies in terms of asymptotic activity.

The other three-cluster networks show similar effects, although the “costs” of overambitious seeding are quantitatively much smaller. For example, in the case of network *klm*, it is clear that the key ingredient in maximizing long-term adoption is to seed cluster *k*. Seeding the other clusters is, at best, superfluous throughout the simulated range and incurs some small costs as p_a grows. In the cases of networks *qrs* and *wxy*, a similar story holds for cluster *q* and *y* respectively. For network



tw , seeding *either* cluster t or cluster u is sufficient, and seeding v is superfluous. In all cases, the zoomed-in views on the right-hand side of Figure 4 show that, at high p_a , a single-cluster-seeding strategy performs best, although as noted above, the “costs” of other seeding strategies are often very small.

7 Simulation Results: k -Core Seeding

We now turn our attention to a different type of seeding that can lead to better long-term outcomes than universal seeding, even when there is no obvious cluster-structure to leverage. Specifically, we will consider seeding, at $t = 0$, only the k -core of the network under consideration. Here, the k -core is defined as usual: it is the subnetwork that remains after repeatedly removing people with degree less than k and all friendship edges incident to these people. The k -core, if it exists, thus corresponds to a dense subnetwork of the original network. Such a seeding approach has been motivated by much of the argumentation above. In particular,

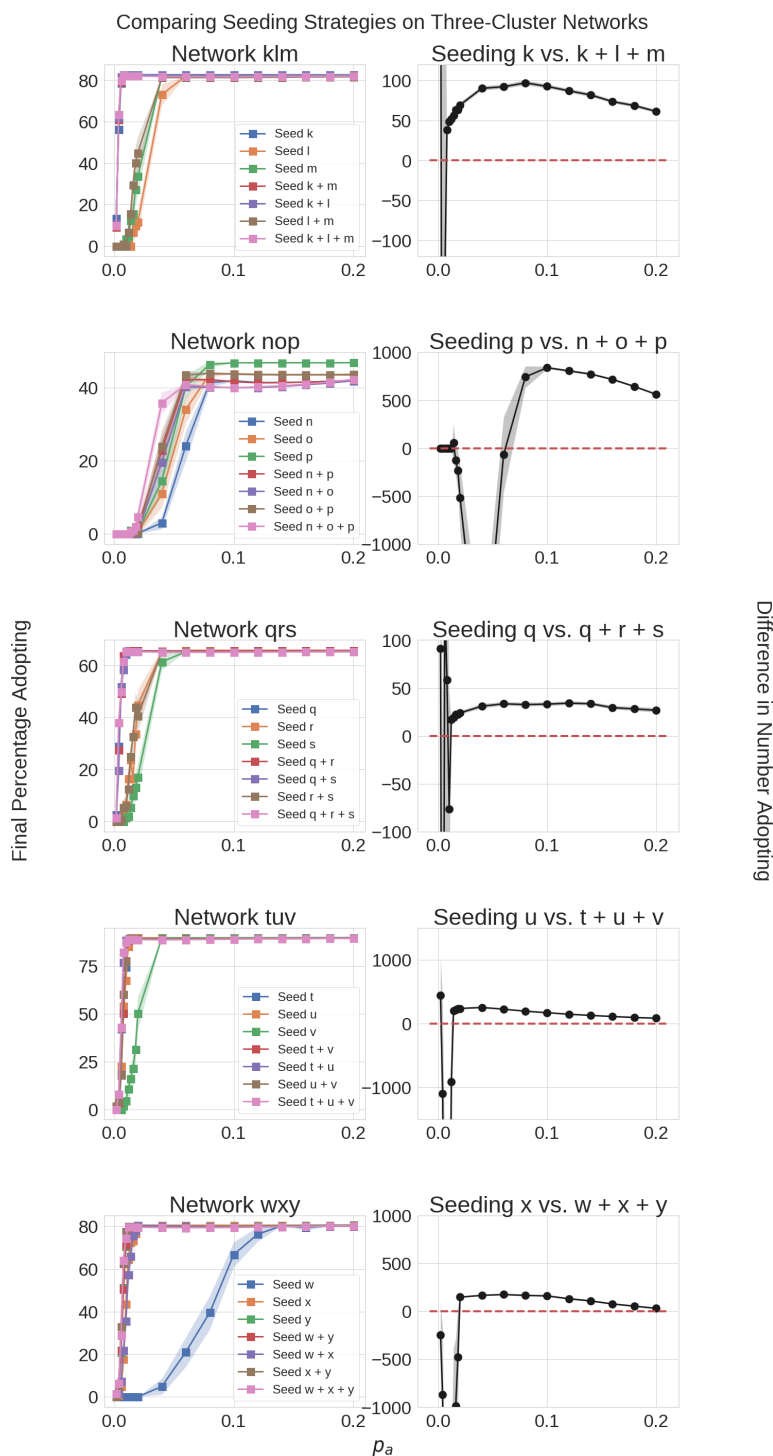


Figure 4 Comparison of cluster-based seeding approaches in the threshold model. Different rows correspond to different three-cluster networks from Table 2. The left-hand column shows the final percentage of the network adopting under various cluster-based seeding strategies; the right-hand column shows the difference in the number adopting under seeding the single cluster with the highest median degree vs. all three clusters. The parameter $s_i = 5$ in these simulations. Each data point is an average over 100 simulations.

it is motivated by the toy examples for both models, where seeding a dense core of the network can be preferable to seeding the entire network.

Repeated-Usage Model: Figure 5 compares seeding the 10-core of various SH clusters to seeding the entire cluster at time $t = 0$ in the repeated-usage model. As in Section 6, we set all $s_i = 2$ here, meaning that everyone needs two active friends to be satisfied during a product experience. We vary p_0 and check which strategy (10-core seeding or universal) wins out in the long-time limit. The asymptotic access and activity values plotted in Figure 5 are again averages over the last 100 time steps of 10000 time-step simulations.

The simulation results of Figure 5 show that 10-core seeding leads to more long-term adoption than universal seeding over large ranges of the low p_0 regime for five different SH clusters. The costs of universal seeding, as compared to the 10-core strategy, are often very large. Our interpretation of these results, echoing our analysis of the toy examples of Section 4, is that it is preferable to allow activity to build up in the core before expanding access to the periphery. This is because people in the core, by virtue of having more friends overall, are much more likely to have satisfying experiences when rates of activity are low. Meanwhile, people in the periphery, because they depend on the usage of a few friends in order to have satisfying experiences, are more likely to have positive product experiences if they receive access after activity has built up in the core.

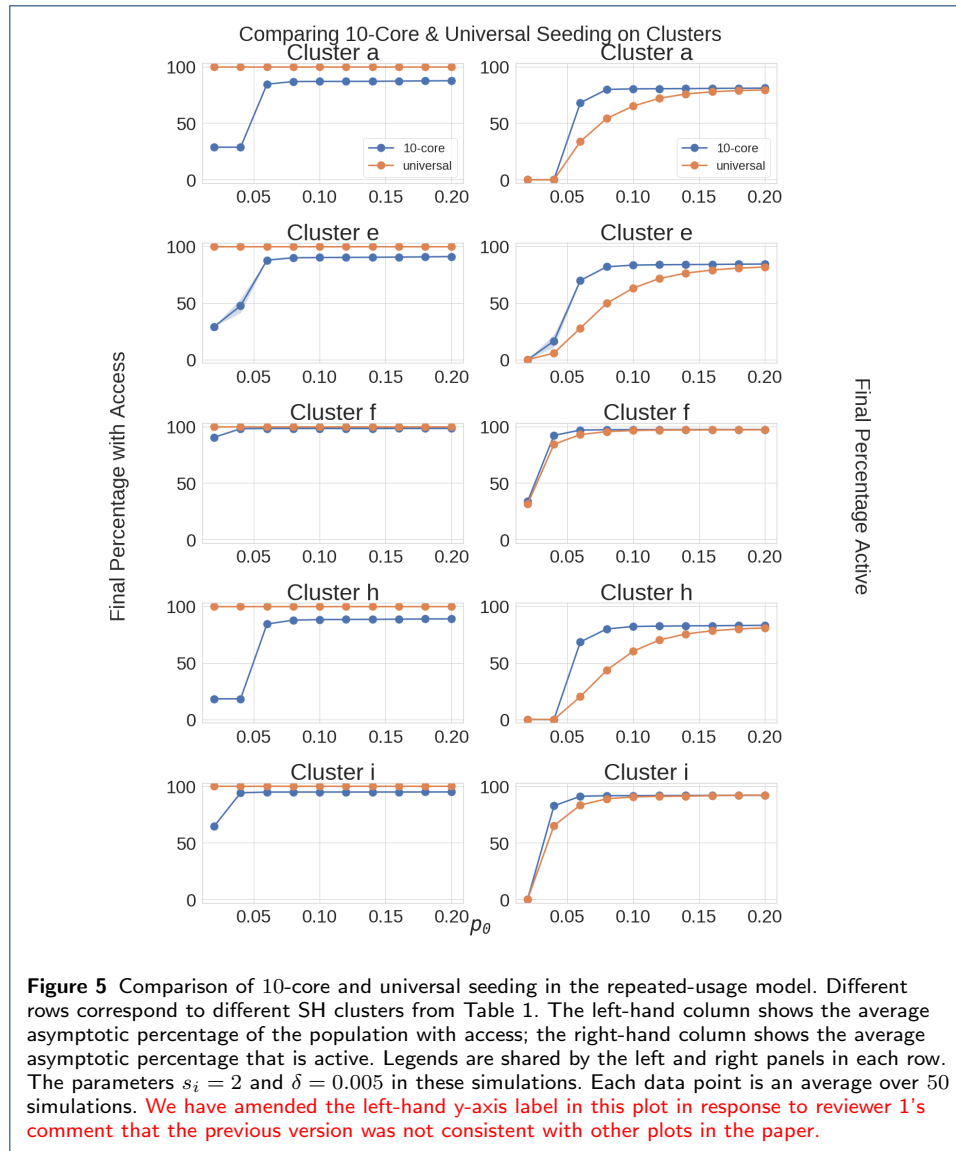
Threshold Model with Churn: We now study k -core seeding in the threshold model. We will again set $s_i = 5$ in these simulations, meaning that each person needs five active friends to become or remain active. We will compare the strategy of seeding the 10-core of each SH cluster to seeding the entire cluster.

The simulation results in Figure 6 generally show three different regimes of behavior. At very low p_a , neither approach leads to substantial long-term adoption. As p_a grows, there is a regime where universal seeding outperforms 10-core seeding. At still higher p_a , 10-core seeding generally wins out. In several cases, 10-core seeding wins by a few percentage points in terms of the total cluster size (clusters a , e , and h). In others, the benefits of 10-core seeding are smaller, but still statistically robust (clusters b , d , f , i , and j).

Again, the tradeoff here is similar to what we have observed previously: in the regime where universal seeding outperforms 10-core seeding, the benefits of early activity in the periphery for the core outweigh the costs to the periphery. Generally though, at high enough p_a , the tradeoff flips, with the costs to the periphery outweighing benefits to the core. Thus, the more conservative seeding strategy (i.e., 10-core seeding) prevails.

8 Discussion: When is Overambitious Seeding Costly?

As we noted above in Section 3, in the repeated-usage model, the question of overambitious seeding amounts to: is it beneficial to seed everyone if everyone whom you seed will accept, but will use at a low rate? On the other hand, in the threshold model, the question is: is it beneficial to seed everyone if only some of those people will accept? Our simulation results show that the answer to both of these questions can be “no” and that various conservative seeding strategies can do substantially better. In this section, we will attempt to abstract from these observations some general principles around *when* overambitious seeding should be a cause for concern.



In both cases, *context* is the key factor in explaining why overambitious seeding is costly. If seeded individuals adopt in contexts where insufficiently many of their friends are adopting or where their friends are not using sufficiently often, they may churn. Meanwhile, if these same individuals are *not* seeded, better contexts may emerge at later times for them to begin using the product. Thus, one rule-of-thumb for *when* to worry about overambitious seeding is the following: overambitious seeding can be costly whenever seeding results in contextually-unaware adoption choices (e.g., people adopting uniformly at random, people using at a rate that's independent of their friends' rates) but where continued usage crucially depends on context.

Note, however, that the effects of overambitious seeding are much more pronounced in the repeated-usage model than in the threshold model. To understand why this is the case, we should note one important distinction between the parameter p_a in the threshold model and the parameter p_0 in the repeated-usage model. In the threshold model, the parameter p_a influences both whether a person adopts the

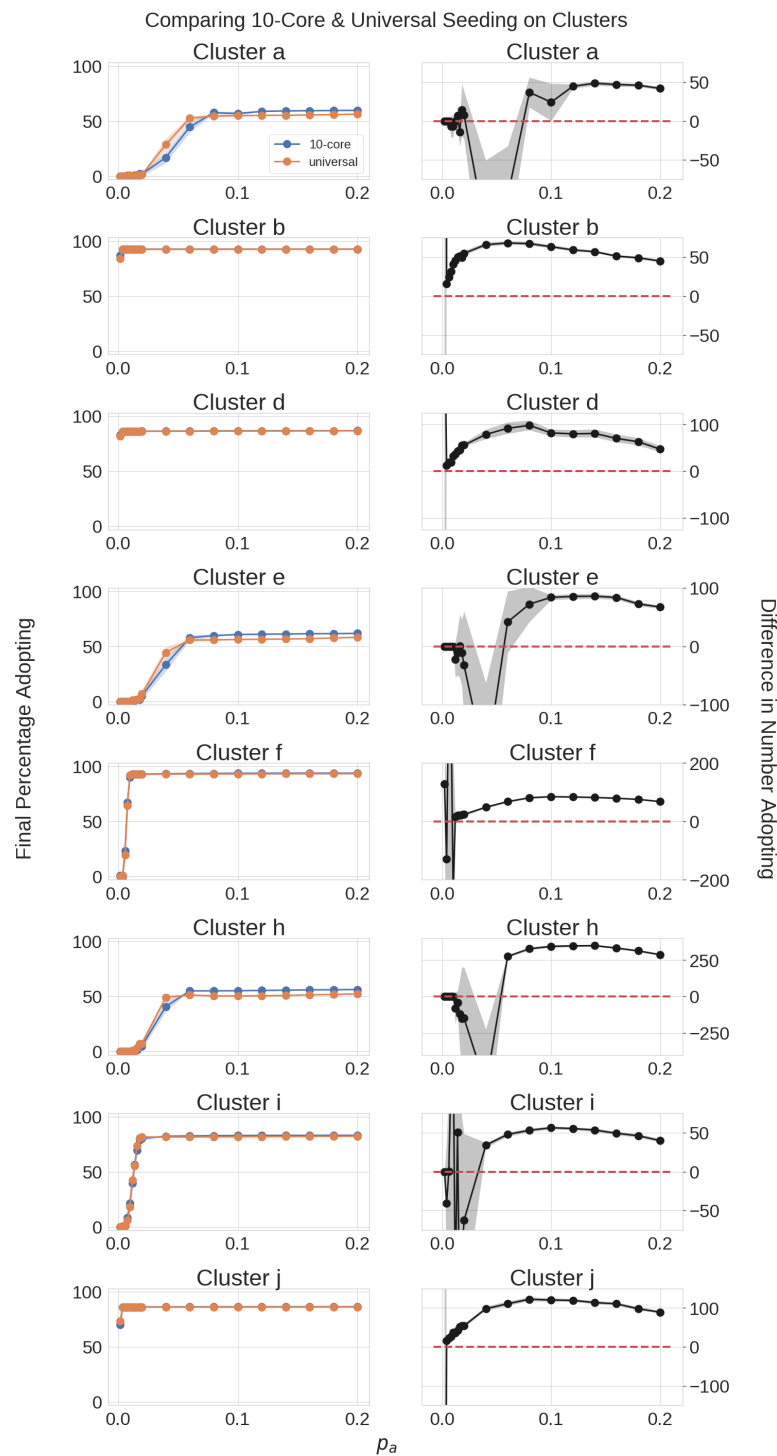


Figure 6 Comparison of 10-core and universal seeding in the threshold model. Different rows correspond to different SH clusters from Table 1. The left-hand column shows the final percentage of the network adopting under various cluster-based seeding strategies; the right-hand column shows the difference in the number adopting under seeding the 10-core vs. the entire cluster. The parameter $s_i = 5$ in these simulations. Each data point is an average over 100 simulations.

product at all *and* how much social support that person can expect at early times. When p_a is low, a person can expect little social support, but it is also less likely that it matters, since the person is also less likely to adopt in the first place. When p_a is higher, a person is more likely to “accept” the product, but is also more likely to experience social contexts that favor continued usage. This restricts the range of p_a where overambitious seeding is likely to be relevant. It also restricts the magnitude of the effect because, typically, the people who incur the costs of overambitious seeding are those who adopt during the seeding round (i.e., in a context-unaware way); p_a constrains this proportion of the population.

We can contrast this with the role of p_0 in the repeated-usage model. Here, p_0 determines how much social support a person can expect at early times and also determines the time scale over which a person will choose to have his or her initial product experiences. Meanwhile, this parameter does *not* determine whether the person has product experiences at all. At low values of p_0 , a large proportion of people can still end up having bad experiences and churning. Hence, there is both a wider range of p_0 where overambitious seeding can be relevant *and* the proportion of the population that can be “lost” due to a bad seeding strategy is large. This suggests another principle around when we should be especially wary of the costs of overambitious seeding: the problem can be especially severe when people’s initial decisions to adopt the product are less correlated with the amount of social support that they can receive at early times.

It is interesting to also consider recent related work by Sela et al. in this context. These authors study product adoption through an SIR model, where a person adopts (transitioning from the S to the I state) either in a budgeted seeding round or because they subsequently have enough adopting friends. After adopting, a person transitions from the influential (I) to non-influential (R) state after a fixed amount of time. When there is a seeding budget b and people are prioritized for seeding by eigenvector centrality, Sela et al. find that the final adoption rate is non-monotonic in the budget b . They call this phenomenon the “flip anomaly” [30]. The “flip anomaly” of Sela et al. also admits a contextual explanation along the lines of those that we have proposed above: if a seeded person is the only adopting friend in a non-seeded person’s local network, then the non-seeded person may not adopt before the seeded person becomes non-influential. If better contexts for the non-seeded person’s adoption emerge later on, the now non-influential friend has no opportunity to contribute to that adoption [30].

There are two interesting points of comparison between the model of Sela et al. and those that we have studied here. First, Sela et al. note that their “flip anomaly” must reverse as the budget grows, because adoption is universal in their model [30]. This is also true of other models with similar properties that have recently been reviewed by Centola [31]. Meanwhile, our results show how analogues of the flip anomaly of Sela et al. can still persist with no seeding budget. Indeed, based on the arguments in this paper, we conjecture that the “flip anomaly” would persist in the unbudgeted case of the model of Sela et al. if adoption under seeding was probabilistic rather than universal. A perhaps more interesting distinction is that Sela et al. show how overambitious seeding can be costly even in the absence of churn, because someone in the R state of their SIR model is still interpreted as

an adopter. This shows that the “possibility of permanent churn” assumption that we encoded into both of the models studied in this paper is not strictly necessary for overambitious seeding to be a problem. Instead, we can make a more general conjecture: overambitious seeding can be costly whenever it results in premature exhaustion of opportunities for further spreading that would better be delayed to later in the spreading process.

9 Conclusion

In this paper, we have revisited a question that we originally posed in [1]: suppose a product developer wants to introduce a new social product to a population of potential adopters and is unconstrained by any seeding budget. In this case, should the developer give the product to everyone immediately (as implied by many classic influence-maximization models), or should the developer adopt a more conservative approach?

We have extended the results of [1] in various ways:

- 1 We have shown that overambitious seeding is not just a concern in the repeated-usage model of [1] but can be a problem in more traditional threshold models as well, once the possibility of churn is introduced.
- 2 We have studied both types of models analytically on certain simplified network structures and thereby developed intuitions for *why* overambitious seeding can be costly.
- 3 We have explored k -core seeding as an alternative to cluster-based seeding, showing that the results of our earlier work are not tied to the cluster-based approach; there are multiple conservative seeding strategies that can outperform seeding everyone at once.

Drawing upon simulation results, we have proposed some general principles around when the possibility of overambitious seeding ought to be considered:

- 1 Overambitious seeding is a concern whenever early adoption can result in the premature exhaustion of a resource for future spreading that would be better delayed to a more favorable context for that spreading.
- 2 Overambitious seeding is *especially* a concern when people’s initial decisions to adopt the product are less correlated with the amount of social support that they can receive at early times.

We emphasize again that the models considered here exclude other pathways to overexposure in the influence maximization problem, including negative word-of-mouth, direct costs to rejection of the seed, and hipster effects. We have excluded these effects to make the case that overambitious seeding can be detrimental in the context of social products, even if none of these factors are at play. Of course, all of these alternative mechanisms are important in real-world settings, and together with the mechanism discussed in this paper, they make the case that product developers should *not* always expend all of the marketing resources at their disposal.

There are many possible interesting extensions of this work. For example, we have always assumed a homogeneous value of p_0 and p_a across the whole population of potential adopters. The costs of overambitious seeding will vary if this assumption of homogeneity is relaxed. If people with many friends have higher values of p_0 or p_a and the friendship network is assortative by degree, presumably seeding everyone

would produce an outcome that is more similar to just seeding the core, diminishing the costs of overambitious seeding. On the other hand, if p_0 and p_a are negatively correlated with degree, that could exacerbate the costs of overambitious seeding and make the considerations of this paper more important.

Another consideration that would mitigate the costs of overambitious seeding in the threshold model would be to allow multiple adoptions per person (i.e., if a person is willing to adopt m times, where $m > 1$). This is because adoptions after the first would happen in contextually aware (and thus, favorable) circumstances, because the person has enough adopting neighbors to adopt. In such a setting, it would be interesting to ask if enriching the model with other aspects of real-world complexity (e.g., some amount of spontaneous churn, some within-person variance over time in social expectations for the product) might reintroduce the costs of overambitious seeding, or fundamentally change the tradeoffs considered here in some unforeseen way.

Here and in [1], we have always taken the perspective that what matters in influence maximization is adoption in the long-time limit. However, it is possible to consider scenarios where there are time constraints, and the goal is to maximize adoption within a fixed time [32]. This too could fundamentally change the tradeoffs discussed here, perhaps shifting them in favor of seeding less conservatively.

On the other hand, a very interesting recent line of work in the influence maximization literature considers other target outcomes besides maximum adoption [33, 34, 35, 36, 37, 38]. As one example, references [33] and [34] consider maximizing the diversity of information shared in a social network. Are “overambitious seeding” considerations relevant in such a setting, or is seeding as widely as possible beneficial for promoting diversity? This seems like a fruitful question to pursue, given the findings of this paper for the more traditional influence-maximization problem.

Funding

~~Not applicable:~~The authors are employed by Facebook and used Facebook computing resources for this research.

Availability of data and materials

The network datasets used in the simulation studies in this paper are not publicly available. However, the simulation code can be made available upon request.

Competing interests

The authors declare that they have no competing interests.

Author's contributions

SI formulated the argument for costs of overambitious seeding in the context of the threshold model, developed and ran the simulations, and wrote the manuscript. LAA proposed the repeated-usage model as a setting where the costs of overambitious seeding might be more pronounced and contributed to the composition of reference [1].

Acknowledgements

We thank Udi Weinsberg, Israel Nir, and Jaroslaw Jankowski for helpful discussions, Shuyang Lin for development of the original simulation infrastructure, and Justin Cheng for reviewing code.

References

1. Iyer, S., Adamic, L.A.: The costs of overambitious seeding of social products. In: International Workshop on Complex Networks and Their Applications, pp. 273–286 (2018). Springer
2. Coleman, J.S., Katz, E., Menzel, H.: Medical Innovation: A Diffusion Study. Bobbs-Merrill Co, New York (1966)
3. Coleman, J., Katz, E., Menzel, H.: The diffusion of an innovation among physicians. *Sociometry* **20**(4), 253–270 (1957)
4. Coleman, J., Menzel, H., Katz, E.: Social processes in physicians' adoption of a new drug. *Journal of chronic diseases* **9**(1), 1–19 (1959)
5. Rogers, E.M.: Diffusion of Innovations. Free Press of Glencoe, New York (1962)

6. Bass, F.M.: A new product growth for model consumer durables. *Management science* **15**(5), 215–227 (1969)
7. Domingos, P., Richardson, M.: Mining the network value of customers. In: *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 57–66 (2001). ACM
8. Hinz, O., Skiera, B., Barrot, C., Becker, J.U.: Seeding strategies for viral marketing: An empirical comparison. *Journal of Marketing* **75**(6), 55–71 (2011)
9. Yadav, A., Wilder, B., Rice, E., Petering, R., Craddock, J., Yoshioka-Maxwell, A., Hemler, M., Onasch-Vera, L., Tambe, M., Woo, D.: Bridging the gap between theory and practice in influence maximization: Raising awareness about hiv among homeless youth. In: *IJCAI*, pp. 5399–5403 (2018)
10. Wilder, B., Onasch-Vera, L., Hudson, J., Luna, J., Wilson, N., Petering, R., Woo, D., Tambe, M., Rice, E.: End-to-end influence maximization in the field. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1414–1422 (2018). International Foundation for Autonomous Agents and Multiagent Systems
11. Banerjee, A., Chandrasekhar, A.G., Duflo, E., Jackson, M.O.: The diffusion of microfinance. *Science* **341**(6144), 1236498 (2013)
12. Goldenberg, J., Libai, B., Muller, E.: Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing letters* **12**(3), 211–223 (2001)
13. Granovetter, M.: Threshold models of collective behavior. *American journal of sociology* **83**(6), 1420–1443 (1978)
14. Schelling, T.C.: *Micromotives and Macrobehavior*. WW Norton, New York (1978)
15. Watts, D.J.: A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences* **99**(9), 5766–5771 (2002)
16. Maximizing the spread of influence through a social network
17. Li, Y., Fan, J., Wang, Y., Tan, K.-L.: Influence maximization on social graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering* (2018)
18. Abebe, R., Adamic, L.A., Kleinberg, J.M.: Mitigating overexposure in viral marketing. In: *Proceedings of the 32nd Conference on Artificial Intelligence* (2018). AAAI
19. Byers, J.W., Mitzenmacher, M., Zervas, G.: The groupon effect on yelp ratings: a root cause analysis. In: *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 248–265 (2012). ACM
20. Kiesling, E., Günther, M., Stummer, C., Wakolbinger, L.M.: Agent-based simulation of innovation diffusion: a review. *Central European Journal of Operations Research* **20**(2), 183–230 (2012)
21. Anderson, E.W.: Customer satisfaction and word of mouth. *Journal of service research* **1**(1), 5–17 (1998)
22. Cui, F., Hu, H.-h., Cui, W.-t., Xie, Y.: Seeding strategies for new product launch: The role of negative word-of-mouth. *PloS one* **13**(11), 0206736 (2018)
23. Alkemade, F., Castaldi, C.: Strategies for the diffusion of innovations on social networks. *Computational Economics* **25**(1-2), 3–23 (2005)
24. Juul, J.S., Porter, M.A.: Hipsters on networks: How a minority group of individuals can lead to an antiestablishment majority. *Physical Review E* **99**(2), 022313 (2019)
25. Jankowski, J., Bródka, P., Kazienko, P., Szymanski, B.K., Michalski, R., Kajdanowicz, T.: Balancing speed and coverage by sequential seeding in complex networks. *Scientific reports* **7**(1), 891 (2017)
26. Adler, J.: Bootstrap percolation. *Physica A: Statistical Mechanics and its Applications* **171**(3), 453–470 (1991)
27. Borgatti, S.P., Everett, M.G.: Models of core/periphery structures. *Social networks* **21**(4), 375–395 (2000)
28. Shalita, A., Karrer, B., Kabiljo, I., Sharma, A., Presta, A., Adcock, A., Killapi, H., Stumm, M.: Social hash: An assignment framework for optimizing distributed systems operations on social networks. In: *NSDI*, pp. 455–468 (2016)
29. Kabiljo, I., Karrer, B., Pundir, M., Pupyrev, S., Shalita, A.: Social hash partitioner: a scalable distributed hypergraph partitioner. *Proceedings of the VLDB Endowment* **10**(11), 1418–1429 (2017)
30. Sela, A., Shmueli, E., Goldenberg, D., Ben-Gal, I.: Why spending more might get you less, dynamic selection of influencers in social networks. In: *Science of Electrical Engineering (ICSEE), IEEE International Conference on The*, pp. 1–4 (2016). IEEE
31. Centola, D.: *How Behavior Spreads: The Science of Complex Contagions vol. 3*. Princeton University Press, Princeton (2018)
32. Chen, W., Lu, W., Zhang, N.: Time-critical influence maximization in social networks with time-delayed diffusion process. In: *Twenty-Sixth AAAI Conference on Artificial Intelligence* (2012)
33. Matakos, A., Gionis, A.: Tell me something my friends do not know: Diversity maximization in social networks. In: *2018 IEEE International Conference on Data Mining (ICDM)*, pp. 327–336 (2018). IEEE
34. Aslay, C., Matakos, A., Galbrun, E., Gionis, A.: Maximizing the diversity of exposure in a social network. In: *2018 IEEE International Conference on Data Mining (ICDM)*, pp. 863–868 (2018). IEEE
35. Tsang, A., Wilder, B., Rice, E., Tambe, M., Zick, Y.: Group-fairness in influence maximization. *arXiv preprint arXiv:1903.00967* (2019)
36. Chen, H., Loukides, G., Fan, J., Chan, H.: Limiting the influence to vulnerable users in social networks: A ratio perspective. In: *International Conference on Advanced Information Networking and Applications*, pp. 1106–1122 (2019). Springer
37. Pasumarthi, R., Narayanam, R., Ravindran, B.: Near optimal strategies for targeted marketing in social networks. In: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pp. 1679–1680 (2015). International Foundation for Autonomous Agents and Multiagent Systems
38. Loukides, G., Gwadera, R.: Preventing the diffusion of information to vulnerable users while preserving pagerank. *International Journal of Data Science and Analytics* **5**(1), 19–39 (2018)