

Strategyproof Mechanisms for Competitive Influence in Networks

Allan Borodin¹ · Mark Braverman² ·
Brendan Lucier³ · Joel Oren¹

Received: 16 December 2013 / Accepted: 20 May 2016 / Published online: 2 June 2016
© Springer Science+Business Media New York 2016

Abstract Motivated by applications to word-of-mouth advertising, we consider a game-theoretic scenario in which competing advertisers want to target initial adopters in a social network. Each advertiser wishes to maximize the resulting cascade of influence, modeled by a general network diffusion process. However, competition between products may adversely impact the rate of adoption for any given firm. The resulting framework gives rise to complex preferences that depend on the specifics of the stochastic diffusion model and the network topology. We study this model from the perspective of a central mechanism, such as a social networking platform, that can optimize seed placement as a service for the advertisers. We ask: given the reported budgets of the competing firms, how should a mechanism choose seeds to maximize overall efficiency? Beyond the algorithmic problem, competition raises issues of strategic behaviour: rational agents should be incentivized to truthfully report their advertising budget. For a general class of influence spread models, we show that when there

✉ Joel Oren
oren@cs.toronto.edu

Allan Borodin
bor@cs.toronto.edu

Mark Braverman
mbraverm@cs.princeton.edu

Brendan Lucier
brlucier@microsoft.com

¹ Department of Computer Science, University of Toronto, 10 King's College Road, Toronto, ON M5S 3G4, Canada

² Department of Computer Science, Princeton University, 35 Olden Street, Princeton, NJ 08540, USA

³ Microsoft Research, 1 Memorial Drive, Cambridge, MA 02142, USA

are two players, the social welfare can be $\frac{e}{e-1}$ -approximated by a polynomial-time strategyproof mechanism. Our mechanism uses a dynamic programming procedure to randomize the order in which advertisers are allocated seeds according to a greedy method. For three or more players, we demonstrate that under an additional assumption (satisfied by many existing models of influence spread) there exists a simpler strategyproof $\frac{e}{e-1}$ -approximation mechanism; notably, this natural greedy mechanism is not necessarily strategyproof when there are only two players.

Keywords Game theory · Social networks · Mechanism design · Influence diffusion

1 Introduction

The concept of word-of-mouth advertising is built upon the idea that referrals between individuals can lead to a contagion of opinion in a population. In this way, a small number of initial adopters can generate a cascade of influence, significantly impacting the adoption of a new product. While this concept has been very well studied in the marketing and sociology literature [1–5], recent popularity of online social networking has made it possible to obtain rich data and directly target individuals based on network topology. Indeed, a potential advantage of advertising served via online social networks is that the platform could preferentially target central individuals, impacting the overall effectiveness of its advertisers' campaigns.

Various models of network influence spread have arisen recently in the literature, with a focus on the algorithmic problem of deciding which individuals to target as initial adopters (or “seeds”) [6–8]. One commonality among many of these (stochastic) models is that the expected number of eventual adopters is a non-decreasing sub-modular function of the seed set. This implies that natural greedy methods [9] can be used to choose initial adopters so as to approximately maximize an advertiser's expected influence. Of course, actually applying such algorithms requires intimate knowledge of the social network, which may not be readily available to all advertisers. However, the owners of the network data (e.g. Facebook or Google) could more easily find potentially influential individuals to target. Our goal is to study the problem faced by a network platform that wishes to provide this service to its advertisers.

Consider the following framework. An online social network platform sells advertising space by contract, offering a price per impression to advertising firms. Each firm has an advertising budget, which determines a number of ad impressions they wish to display. As an additional service to the firms, the platform attempts to optimize the placement of advertisements so to maximize influence diffusion. This optimization is to be provided as a service to the advertisers, with the primary goal of making the social network more attractive as a marketing platform. The network provider thus faces an algorithmic problem: maximize the total influence of the advertisers given their budgets (i.e. number of impressions). This problem may be complicated by competition between advertisers, which results in negative externalities upon each others' product adoptions. Moreover, since advertising budgets are private, there is also a game-theoretic component to the problem: the placement algorithm should not

incentivize firms to reduce their budgets. This may happen if, due to properties of the algorithm, lower-budget advertisers might obtain higher expected influence than advertisers with higher budgets.

Crucial to this problem formulation is the way in which influence is modeled by the advertising platform. We present a general submodular assignment problem with negative externalities, which captures most previous influence models that have been proposed in the literature [10–13]. Within this framework, we consider the optimization problem faced by a central mechanism that must determine the seed nodes for *each* advertiser, given the advertisers' budget constraints. The goal of the mechanism is to maximize the overall efficiency of the marketing campaigns, but the advertisers are strategic and may under-report their budgets to increase their own product adoption rates.

Two points of clarification are in order. First, our formulation differs from a line of prior work that studies equilibria of the game in which each advertiser selects their seed set directly [11, 13]. Such a game supposes that each advertiser has detailed knowledge of the social network topology, the ability to compute or converge to equilibrium strategies, and the power to target arbitrary individuals in the network. Our work differs in that we assume that the targeted advertising goes through an intermediary (the social network), which selects seed sets on the players' behalf.

Second, we suppose that advertiser budgets and the price per impression are set exogenously (or, alternatively, that the seeds correspond to special offers or other interventions of limited quantity). As such, we do not explicitly model the problem of maximizing revenue; rather, the role of our mechanism is to decide where to place the purchased impressions so as to maximize the social welfare (i.e. the total influence spread). In this sense our framework is closer in spirit to matching algorithms for display advertising [14, 15] than to revenue-optimal mechanism design. There are many ways in which this model could be enriched, such as by endogenizing budgets or allowing complex pricing schemes that depend upon expected influence. We leave these as avenues for future work, though we note that such extensions presuppose that agents have sufficient knowledge of the spread process and graph topology to accurately value initial adoption sets.

Our general model of competitive influence captures prior models from the literature such as those discussed in the related works Sect. 1.2 to follow. (Further discussion of the relation to other models appears in Appendix “Relation with Other Diffusion Models”.) The model treats influence spread as a submodular process, with some extra conditions that relate the influence of different firms; namely, an adverse competition assumption that reflects the competitive nature of the process, as well as mechanism indifference and agent indifference assumptions¹ that will be introduced in Sect. 1.1. Our model of competitive influence spread is described formally in Sect. 2.

¹ In the conference version of this work [16], we stated a strategyproof 2-approximation mechanism for the 2-player case that allowed for more general spread processes. That argument had a flaw and we are now restricting our results to spread processes that satisfy the mechanism indifference property. Our mechanism for three or more players uses both the mechanism and agent indifference properties, as in the conference version.

We wish to design mechanisms that are strategyproof, in that rational agents are incentivized to truthfully reveal their budgets. In particular, an agent should not be able to increase its expected influence by reducing its requested number of seeds (i.e. budget). The difficulty in avoiding such non-monotonicities is that the expected outcome of an advertiser can be negatively impacted by externalities imposed by the allocation to its opponents, which can depend on the budget declarations in a non-trivial manner.

1.1 Our Results

We design two different $\frac{e}{e-1}$ -approximate strategyproof mechanisms for the competitive influence maximization problem. Our main result is the strategyproof mechanism for use when there are two competing advertisers. This result assumes that the social welfare (i.e. the expected number or weight of nodes that are eventually influenced) is independent of the manner in which elements are partitioned among the players (*mechanism indifference*). This mechanism uses a novel technique for monotonicizing the expected utilities of the agents using geometric properties of the problem in the two-player case.

Our second and more natural mechanism is for three or more players, under an additional (and also common) natural restriction on the influence spread process. Specifically, we require that the payoff of a player does not depend on the manner in which the elements allocated to her competitors are partitioned among the competitors (*agent indifference*). We note that both the mechanism and agent indifference assumptions are implicit in many prior models of influence spread [10], [11]. Interestingly, our analysis here makes crucial use of the presence of three or more players, and indeed we show in Sect. 3.1 that this mechanism fails to be strategyproof when only two players are present, even with these two additional assumptions.²

Our mechanisms run in time polynomial in the budgets submitted by the agents and in the size of the underlying ground set. This dependence on the budget values is necessary, as the mechanism constructs a solution consisting of sets of this size. Our dependence on the size of the underlying ground set is captured by oracle queries for an element that maximizes a marginal increase in social welfare. Given oracle access to queries of this nature, our algorithm would run in time polynomial in the declared budgets. Generally speaking, the spread process itself is randomized and as in [6, 7], the oracle can be viewed as providing an element that approximately maximizes the marginal gain by sampling enough trials of this process [6, 7]. Our analysis also holds when such approximate marginal maximizers are used to implement our underlying greedy algorithm; following the exposition in [17], such an approximate maximizer provides an approximation that approaches $\frac{e}{e-1}$ as the oracle approximation approaches 1. We will simplify our discussion throughout by assuming it is possible to find elements that exactly maximize marginal gains in social welfare.

² Notice that the agent indifference property holds vacuously in the two-player case, as there is only one other player

1.2 Related Work

Models of influence spread in networks, covering both cascade and threshold phenomena, are well-studied in the sociology and marketing literature [1,2]. The (non-competitive) problem of maximizing influence in social networks was theoretically modelled by Kempe et al. [6,7]. Subsequent papers extended these models to a competitive setting in which there are multiple advertisers. Carnes et al. [10] suggested the Wave Propagation model and the Distance Based model, which were based on the Independent Cascade model. Additionally, Dubey et al. [18], Bharathi et al. [11], Kosta et al. [19], and Apt et al. [20] also studied various competitive models. The main issue that these models addressed was how to arbitrate ties in each step of the process, determining which technology a node will assume when reached by several technologies at once. The main algorithmic task addressed by these models is choosing the optimal set of nodes for a player entering an existing market, in which the competitor's choice of initial nodes is already known. Borodin et al. [12] presented the OR model which proposes a different approach, in which the previously studied, non-competitive diffusion models proceed independently for each technology as a first phase of the process, after which the nodes decide between each technology according to some decision function.

Independent of our initial work [21], Goyal and Kearns [13] provided bounds on the efficiency of equilibria (i.e. the price of anarchy) in a competitive influence game played by two players. Their influence spread model is characterized by switching functions (specifying the process by which a node decides to adopt a product) and selection functions (specifying the manner in which nodes decide which product to adopt). They demonstrate that an equilibrium of the resulting game yields half of the optimal social welfare, given that the switching functions are concave. Their model is closely related to our own. Specifically, the social welfare function is monotone and satisfies the mechanism indifference assumption, and concavity of the switching function implies that the social welfare is submodular (by [8]), so our mechanism for two players applies to their model as well³. Goyal and Kearns also note that their results extend to $k > 2$ players, resulting in an approximation factor of $2k$, when the selection function is linear; this linearity implies our agent indifference assumption, and hence our mechanism for three or more players also applies. However, we note that the Goyal and Kearns results on efficiency at equilibrium are satisfied without an intervening mechanism and hence are incomparable with the mechanism results of this paper.

Finally, to the best of our knowledge, there is only one other paper that considers a mechanism design problem in the context of competitive influence spread. Namely, Singer [22] considers a social network where the nodes are viewed as agents who have private costs for hosting a product and the mechanism has a budget for inducing some set of initial nodes to become hosts. The mechanism wishes to maximize the number of nodes that will eventually be influenced and each agent wishes to maximize their profit equal to the inducement received minus its private cost.

³ An “adverse competition” assumption in [13] is stated for $k = 2$ agents and holds at every node. Their assumption is somewhat stronger than ours, which we only apply to the social welfare function. See Sect. 2.

2 Preliminaries

We consider a setting in which there is a ground set $U = \{e_1, \dots, e_n\}$ of n elements (e.g. nodes in a social network), and k players. An allocation is some $(S_1, \dots, S_k) \in 2^U \times \dots \times 2^U$; that is, an assignment of set⁴ S_i to each player i .

We are given functions $f_i: 2^U \times \dots \times 2^U \rightarrow \mathbb{R}_{\geq 0}$, denoting the expected values of players $i = 1, \dots, k$, for allocation (S_1, \dots, S_k) . We define $f = \sum_{i=1}^k f_i$, so that $f(\mathbf{S}) = f(S_i, \mathbf{S}_{-i})$ denotes the total expected welfare of the allocation $(\mathbf{S}) = (S_1, \dots, S_k) = (S_i, \mathbf{S}_{-i})$.

We will require that functions f , and f_1, \dots, f_k satisfy certain properties, motivated by known properties of influence spread models studied in the literature. First, we will assume that f is a submodular non-decreasing function, in the following sense. For any $S_i \subseteq S'_i$, \mathbf{S}_{-i} , and $e \in U$, we have $f(S_i, \mathbf{S}_{-i}) \leq f(S'_i, \mathbf{S}_{-i})$ and

$$f(S_i \cup \{e\}, \mathbf{S}_{-i}) - f(S_i, \mathbf{S}_{-i}) \geq f(S'_i \cup \{e\}, \mathbf{S}_{-i}) - f(S'_i, \mathbf{S}_{-i}).$$

We will also require that for all $i = 1, \dots, k$, the function f_i be non-decreasing in the allocation to player i , so that $f_i(S_i, \mathbf{S}_{-i}) \leq f_i(S'_i, \mathbf{S}_{-i})$ for any $S_i \subseteq S'_i$. These assumptions are all rather standard in spread models.

A more substantive property of the social welfare function f that we assume is *mechanism indifference* (MeI) which we define as follows: $f(\mathbf{S}) = f(\mathbf{S}')$ whenever the sets $\bigcup_i S_i$ and $\bigcup_i S'_i$ are equal. In words, this assumption means the social welfare function is invariant, regardless of the manner in which we divide the total set of allocated items. Going back to our motivating setting of competitive influence diffusion processes, this assumption pertains to scenarios in which the final decision of whether to adopt some technology is independent of the division of the total marketing resources (i.e., the initial allocated nodes). We note that to the best of our knowledge, most of the competitive diffusion models that are studied in the literature have this property (see Appendix “Relation with Other Diffusion Models”).

For our result concerning three or more players we will need the *agent indifference* (AgI) property defined as follows: $f_i(S_i, \mathbf{S}_{-i}) = f_i(S_i, \mathbf{S}'_{-i})$ whenever sets $\bigcup_{j \neq i} S_j$ and $\bigcup_{j \neq i} S'_j$ are equal. That is, each agent’s utility depends on the set of items allocated to the other players, but not on how the items are partitioned among those players. We will call this the *Agent Indifference* (AgI) assumption. Notice that in the two-players case, this assumption is vacuous. We again note that many of the existing competitive influence models satisfy this property.

We impose one final model assumption with regard to the competitive aspect of the model, which we call *adverse competition*; namely, each f_i is non-increasing in the allocation to other players. That is, for all $j \neq i$, $f_i(S_j, \mathbf{S}_{-j}) \geq f_i(S'_j, \mathbf{S}_{-j})$ for any $S_j \subseteq S'_j$. This assumption captures our intuition that, in a competitive influence model, the presence of additional adopters for one player can only impede the spread

⁴ For notational convenience we will assume that S_1, \dots, S_k are sets, but our results extend to permit multisets (i.e., where the same element can be awarded multiple times to one agent). See Appendix “Relation with Other Diffusion Models” for further discussion.

of influence for another player. We discuss the motivation for and necessity of this assumption in Appendix “Relation with Other Diffusion Models”.

2.1 The Algorithmic Problem and Overall Approach

We study the following task. We assume an influence spread maximization problem with unit cost to influence an initial node. Given the agents’ budgets $b_1, \dots, b_k \geq 0$, which then denotes the number of initial nodes that each agent wants to initially influence (their reported budgets), we wish to find sets $S_1, \dots, S_k \subseteq V$, with $|S_i| = b_i$, for all $i = 1, \dots, k$, such that $f(S_1, \dots, S_k)$ is maximized. In order to decouple the task of estimating the stochastic influence function from the combinatorial problem at hand, we assume we are given oracle access to the functions f and f_1, \dots, f_k .⁵

Suppose that \mathcal{A} is a deterministic algorithm for the above problem, so that $\mathcal{A}(b_1, \dots, b_k)$ denotes an allocation for any $b_1, \dots, b_k \geq 0$. We say that \mathcal{A} is *monotone* if, for all budget vectors $\mathbf{b} = (b_1, \dots, b_k) \in \mathbb{Z}_{\geq 0}^k$, $f_i(\mathcal{A}(b_i, \mathbf{b}_{-i})) \leq f_i(\mathcal{A}(b_i + 1, \mathbf{b}_{-i}))$, for each player $i = 1, \dots, k$. We extend this definition to randomized algorithms in the natural way, by taking expectations over the outcomes returned by \mathcal{A} .

We will assume that each agent i has a *type* \tilde{b}_i , representing the maximum number of elements they can be allocated. This corresponds to the agent’s true budget. The utility of player i for allocation $\mathbf{S} = (S_1, \dots, S_k)$ is

$$u_i(\mathbf{S}) = \begin{cases} f_i(S_i, \mathbf{S}_{-i}) & \text{if } |S_i| \leq \tilde{b}_i \\ -\infty & \text{otherwise.} \end{cases}$$

Note that the second part of the definition of the utility function essentially means that the agent should not be allocated additional nodes beyond what is permitted by his budget (i.e., the agents have hard budgets). We then say that algorithm \mathcal{A} is *strategyproof* if, for all $\mathbf{b} \in \mathbb{Z}_{\geq 0}^k$ and $b'_i \leq b_i$, $u_i(\mathcal{A}(b'_i, \mathbf{b}_{-i})) \leq u_i(\mathcal{A}(b_i, \mathbf{b}_{-i}))$. In other words, an algorithm is strategyproof if it incentivizes each agent to report its type truthfully.

Due to the above assumptions, we observe that the optimization problem of maximizing the social welfare function $f(\cdot)$ subject to the reported budgets can be stated in the framework of maximizing a submodular set-function subject to a cardinality constraint that has been studied by Nemhauser et al. [23]. Namely, given a non-decreasing and submodular set function $g : 2^X \rightarrow \mathbb{R}_{\geq 0}$ over a universe X , and a cardinality constraint T , the goal is to find a set $S \subseteq X$ that maximizes $g(S)$.

Put in terms of our setting, we can set the cardinality constraint to be $B = \sum_{i=1}^k b_i$, the social welfare function can be re-expressed as the following single-parameter, non-decreasing, and submodular objective function $g(\cdot)$:

⁵ This assumption is compliant with most of the models studied in the literature, in which these values (the spread functions of each technology), can be estimated with arbitrary precision via sampling.

$$g(T) = f(\mathbf{S}), \text{ where } T = \bigcup_{i=1}^k S_i. \quad (1)$$

This alternative objective function $g(T)$ conforms to the single-parameter convention of submodular set-functions. However, we will mostly refer to the former objective $f(\mathbf{S})$ for clarity and succinctness

As a result of this correspondence with the framework of submodular function maximization, we will make use of the well-known greedy algorithm due to Nemhauser et al. [17,23] for selecting the total set of nodes to be allocated to the agents, which will give us an $e/(e-1)$ approximation ratio of the optimal social welfare. The algorithm operates as follows: starting from an empty set $S = \emptyset$, while $|S| < B$, the algorithm adds to S the element $x \in U \setminus S$ that maximizes the marginal contribution to the function $g(\cdot)$. As Nemhauser et al. show, this “hill-climbing” algorithm admits a $e/(e-1)$ approximation to the optimal social welfare.

The hill-climbing algorithm admits the total set of elements that will be allocated to *all* the agents. Thus, the algorithmic problem that we will address is how to partition this total set of nodes among the agents in a manner that incentivizes the agents to truthfully report their budgets. More specifically, if we denote the collective budget of the agents by $B = \sum_{i=1}^k b_i$, we will employ the greedy hill-climbing algorithm to generate the sequence (e_1^*, \dots, e_B^*) of nodes that are to be allocated to the agents, where the ordering is taken with respect to the iterative addition process. Then, our algorithms will decide, for each $j = 1, \dots, B$, to which agent v_j^* will be assigned.

Importantly, notice that due to our MeI assumption, our mechanisms will always allocate the agents *disjoint* sets of vertices, as allocating different agents the same vertex will have no effect on the social welfare. Henceforth, we will assume that our general influence spread model satisfies both the MeI and adverse competition assumptions. We will only need to utilize the AgI assumption in Sect. 4 when we consider the mechanism for $k > 2$ players.

3 A Strategyproof Mechanism for Two Players

In this section we describe our mechanism for allocating nodes when there are two agents. The case of $k > 2$ agents is handled in Sect. 4, under the additional AgI assumption that is vacuously true for the case $k = 2$. Our mechanism is based on the greedy algorithm described in Sect. 2.

3.1 Counter Examples When There are Two Agents

In Sect. 3.2 we carefully construct a distribution for determining which player will receive each of the initial nodes that have been chosen by the natural greedy algorithm. This distribution will induce strategyproofness for two players. To motivate these algorithmic gymnastics, we now demonstrate that more natural orderings fail to result in strategyproofness.




Figure 1 The counter-example for the dictatorship ordering

We begin by considering the “dictatorship” ordering, in which one player is first allocated nodes up to his budget, and only then is the other player allocated nodes. We will refer to the agents as A and B , and their utilities as f_A and f_B respectively; suppose that A is the dictator. For the purposes of our example we will describe f_A and f_B in terms of the following concrete (but simple) competitive influence spread process⁶ on an undirected network $G = (V, E)$.

Suppose that each agent is given an initial seed set, say S_A and S_B . For agent A , each node in S_A is given a single chance to *activate* each of its neighbors independently, which it does with probability $p = 0.9$. (Note that this activation process is not recursive; it affects only the neighbors of S_A). We then, independently, allow each node in seed set S_B to attempt to activate each of its neighbors, resulting in a set of nodes activated by B . To determine the final influence sets, any node activated only by A is influenced by A , any activated only by B is influenced by B , and any node activated by both will choose between the two agents uniformly at random. The value of $f_A(S_A, S_B)$ is the expected number of nodes influenced by A at the end of this process, and similarly for f_B . One can easily show that an agent’s influence is non-decreasing in its seed set, that the sum of influences is submodular non-decreasing, and that the functions satisfy adverse competition.

Our network is as follows (see Fig. 1). The graph consists of two components; one is the complete bipartite graph $K_{2,10}$, and the other is the star $K_{1,4}$. Let w_1 and w_2 be the two nodes of degree 10, and let v be the center of the star. We claim that the greedy algorithm paired with the dictatorship ordering is not strategyproof for this network. Suppose each agent declares a budget of 1; in this case, the algorithm will allocate w_1 to agent A , then it will allocate v to agent B (since $4p > 10(1 - (1 - p)^2) - 10p$, which means that v maximizes the marginal gain in *social welfare*). This results in an

⁶ This process is a simplification of the OR model [12].

expected influence of $10p = 9$ for agent A . In the case where A has a budget of 2 (and B 's budget is still 1), the greedy algorithm will allocate w_1 and v to agent A (for the same reason as before), and will give w_2 to agent B . In this case, the influence of agent A becomes $4p + 10(p \cdot (1 - p) + \frac{p^2}{2}) = 8.55 < 9$, so in particular his influence is not non-decreasing in his declaration.

The above construction can be modified to show that various other agent orderings for the greedy algorithm fail to result in strategyproof mechanisms. Appendix “Counterexamples When There are Two Agents (Extended Discussion)” provides the following examples:

1. The Round Robin ordering: the mechanism alternates between the players when allocating a node.
2. Always choosing the player having the smallest current unsatisfied budget breaking ties in favor of player A .
3. Taking a uniformly random choice over all orderings with the required number of allocations to A and B .

The last example is particularly relevant, since in Sect. 4 we show that for the case of $k > 2$ agents, taking a uniformly random permutation over the allocation in combination with the greedy hill-climbing algorithm results in a strategyproof algorithm that admits an $\frac{e}{e-1}$ approximation to the optimal social welfare. In contrast, for the case of $k = 2$, and even under the extra AgI assumption placed in Sect. 4, the uniformly random mechanism is not strategyproof.⁷

3.2 A Strategyproof Mechanism for Two Agents

Theorem 1 *In the case of two strategic agents, there exists a strategyproof mechanism that admits a $e/(e - 1)$ -approximation of the maximal social welfare. Furthermore, the algorithm runs in polynomial time.*

As previously mentioned, our mechanism is divided into two stages: first, we select the sets of nodes to be allocated to the agents (using the greedy hill-climbing algorithm described in the preliminaries). Then, we randomly construct a permutation that maps the nodes selected in the first stages to the agents in a manner that respects their budgets, and that incentivizes truthful reporting of budgets.

The idea behind our construction, at a high level, is as follows. We will construct the distribution for use with budgets (a, b) recursively. Writing $t = a + b$, we first generate the two distributions for the case $t = 1$ (which are trivial), followed by $t = 2$, etc. To construct the distribution for budgets (a, b) , we consider the following thought experiment. We will choose an ordering in one of two ways. Either we choose a permutation according to the distribution for budget pair $(a - 1, b)$ and then append a final allocation to A , or else choose a permutation according to the distribution for budget pair $(a, b - 1)$ and append an allocation to agent B . If we choose the former

⁷ One can verify that the influence model described above, used for our counterexample, does satisfy both MeI and AgI, although the AgI condition is vacuous.

option with some probability α , and the latter with probability $1 - \alpha$, this defines a probability distribution for budget pair (a, b) .

What we will show is that, assuming our distributions are constructed to adhere to certain invariants, we can choose this α such that the resulting randomized algorithm (i.e., the greedy algorithm applied to permutations drawn from the constructed distributions) will be monotone. That is, the expected influence of agent A under the distribution for (a, b) is at least that of the distribution for $(a - 1, b)$, and similarly for agent B . The existence of such an α is not guaranteed in general; we will need to prove that our constructed distributions satisfy an additional “cross-monotonicity” property in order to guarantee that such an α exists.

One problem with the above technique is that it does not bound the size of the support of the distributions. In general there will be exponentially many possible permutations to randomize over, leading to exponential computational complexity to compute the value of α for each pair of a and b . One might attempt to overcome such issues by sampling to estimate the required probabilities, but this introduces the possibility of non-monotonicities due to sampling error, which we would like to avoid. We demonstrate that each distribution we construct can be “pruned” so that its support contains at most three permutations, while still retaining its monotonicity properties. In this way, we guarantee that our recursive process requires only polynomially many queries (to the influence functions) in order to choose a permutation.

3.2.1 The Allocation Algorithm

Our algorithm will proceed by first selecting the set of $a + b$ nodes, using the greedy hill-climbing algorithm, while remembering the order of their selections. The second stage of our algorithm consists of choosing a distribution over orders in which nodes are allocated to the two agents. This will be stored in a matrix M , where $M[a, b]$ contains a distribution over sequences $(y_1, \dots, y_t) \in \{A, B\}^{a+b}$, containing a ‘A’s and b ‘B’s. We then choose a sequence from distribution $M[a, b]$ and greedily construct a final allocation with respect to that ordering. We begin by describing the manner in which the allocation is made, given the distribution over orderings. The algorithm is given as Algorithm 1. An important property of the allocation algorithm that we will require for our analysis is that, given a sequence drawn from distribution $M[a, b]$, the allocation is chosen myopically (due to the nature of the greedy algorithm). That is, items are chosen for the agents in the order dictated by the given sequence, independent of subsequent allocations. We will use this property to construct the distribution $M[a, b]$, which will be tailored to the greedy algorithm to ensure strategyproofness.

Recall that the approximation ratio for the greedy algorithm guarantees that Algorithm 1 will be a $e/(e - 1)$ -approximation to the optimal total influence, and because of the mechanism indifference property, this approximation guarantee will hold regardless of how we partition the total set among the agents. It remains only to demonstrate that we can construct our distributions in such a way that the expected payoff to each agent is monotone non-decreasing in his budget.

Algorithm 1: The Allocation Mechanism for Two Agents**Input:** A ground set $U = \{e_1, \dots, e_n\}$, budgets a, b for agents A and B , respectively**Output:** An allocation $I_A, I_B \subseteq U$ for the two agents

```

1  $S \leftarrow \emptyset$ 
2 for  $i = 1 \dots a + b$  do
3   Let  $e_i^* = \arg \max_{e \in U} (f(S + e) - f(S))$ 
4    $S \leftarrow S \cup \{e_i^*\}$ 
5 end
   /*Build permutation table. */
6 Let  $(e_1^*, \dots, e_{a+b}^*)$  be the ordering of the elements in  $S$ , based on the order of their additions.
7  $S_A \leftarrow \emptyset$ ;  $S_B \leftarrow \emptyset$ 
8  $M \leftarrow \text{ConstructDistributions}(a, b)$ ;
   /* $M[a, b]$  will be a distribution over sequences  $(y_1, \dots, y_{a+b}) \in \{A, B\}^{a+b}$  */
9 Choose  $(y_1, \dots, y_{a+b})$  from distribution  $M[a, b]$ ;
10 for  $i = 1 \dots a + b$  do
11   if  $y_i = 'A'$  then
12      $S_A \leftarrow S_A \cup \{e_i^*\}$ ;
13   else
14      $S_B \leftarrow S_B \cup \{e_i^*\}$ ;
15   end
16 end

```

3.3 Constructing Matrix M

We describe the procedure *ConstructDistributions*, used in Algorithm 1, to generate distributions over orderings of assignments to agents A and B . We will build table $M[\cdot, \cdot]$ recursively, where $M[a, b]$ describes the distribution corresponding to budgets a and b . Our procedure will terminate when the required entry has been constructed.

We think of $M[a, b]$ as a distribution over sequences of the form (y_1, \dots, y_{a+b}) , where $y_i \in \{A, B\}$. For any given sequence, the corresponding allocation is determined since the greedy algorithm applied in Algorithm 1 is deterministic. We can therefore also think of $M[a, b]$ as a distribution over allocations, and in what follows we will refer to “allocations drawn from $M[a, b]$ ” without further comment.

Note that $M[0, b]$ must assign probability 1 to the sequence (B, B, \dots, B) , and similarly $M[a, 0]$ assigns probability 1 to (A, A, \dots, A) . We will construct the remaining entries of the table $M[a, b]$ in increasing order of $a + b$.

Before describing the recursive procedure for filling the table, we provide some notation. Given M , we will write $w^A(a, b)$ for the expected value of agent A under the distribution of allocations returned by $M[a, b]$. Similarly, $w^B(a, b)$ will be the expected value of agent B , and $w(a, b) = w^A(a, b) + w^B(a, b)$ is the expected total welfare. For notational convenience, set $w^A(a, b) = w^B(a, b) = 0$ if $a < 0$ or $b < 0$.

We will construct M so that the following invariants hold for all $a > 0$ and $b > 0$:

1. $w^A(a, b) \geq w^A(a - 1, b + 1)$.
2. $w^A(a, b) \geq w^A(a - 1, b)$.
3. $w^B(a, b) \geq w^B(a, b - 1)$.
4. The support of $M[a, b]$ contains at most 3 sequences.

The first invariant is a type of cross-monotonicity property, which will help us to construct the entries of matrix M . The second two desiderata capture the monotonicity properties we require of our algorithm. Note that if M satisfies these properties, then Algorithm 1 will be monotone and hence strategyproof. The final property limits the complexity of constructing and sampling from $M[a, b]$, implying that Algorithm 1 runs in polynomial time.

We now describe the way in which we construct distribution $M[a, b]$, given distributions $M[a', b']$ for all $a' + b' < a + b$. We consider two distributions: the first selects a sequence according to $M[a - 1, b]$ and appends an “A”, and the second selects a sequence according to $M[a, b - 1]$ and appends a “B”. Call these two distributions D_1 and D_2 , respectively. What we would like to do is find some $\alpha, 0 \leq \alpha \leq 1$, such that if we choose from distribution D_1 with probability α and distribution D_2 with probability $1 - \alpha$, then the resulting combined distribution (for $M[a, b]$) will satisfy $w^A(a, b) \geq w^A(a - 1, b)$ and $w^B(a, b) \geq w^B(a, b - 1)$. Of course, this combined distribution may have support of size up to 6 (3 from D_1 and 3 from D_2) but we will show that it can be pruned to a distribution with the same expected influence for agents A and B, with at most 3 permutations in its support.

Our main technical lemma, Lemma 1, demonstrates that an appropriate value of α , as described in the process sketched above, is guaranteed to exist and can be found efficiently. Before stating the lemma we introduce some helpful notation. Write $\Delta^{\oplus B}(a, b) = w(a, b) - w(a, b - 1) \geq 0$. That is, $\Delta^{\oplus B}(a, b)$ is the marginal gain in total welfare when agent B increases his budget from $b - 1$ to b . Note that due to the manner in which the algorithm operates, $\Delta^{\oplus B}(a, b) \geq 0$, as this is just the marginal increase obtained when adding an item in iteration $a + b$ of the greedy part of the algorithm (lines 3–4). Also note that due to MeI and the manner by which the algorithm selects the initial nodes, this value is solely determined by the sum $a + b$.

Lemma 1 *It is possible to construct table M in such a way that the following properties hold for all $a + b \geq 1$:*

1. $w^A(a, b) \geq w^A(a - 1, b + 1)$
2. $w^A(a, b) \geq w^A(a - 1, b)$
3. $w^A(a, b) \leq w^A(a, b - 1) + \Delta^{\oplus B}(a, b)$

Furthermore, the entries of M can be computed in polynomial time.

Notice that condition 3 in Lemma 1 implies that agent B’s valuation is monotone increasing with his budget:

$$\begin{aligned}
 w^A(a, b - 1) &\geq w^A(a, b) - \Delta^{\oplus B}(a, b) \\
 &= w^A(a, b) - [w(a, b) - w(a, b - 1)] \\
 &= w^A(a, b) - \left[\left(w^A(a, b) + w^B(a, b) \right) \right. \\
 &\quad \left. - \left(w^A(a, b - 1) + w^B(a, b - 1) \right) \right] \\
 &= w^A(a, b - 1) + w^B(a, b - 1) - w^B(a, b) \\
 &\Rightarrow w^B(a, b) \geq w^B(a, b - 1)
 \end{aligned} \tag{2}$$

Proof We will proceed by induction on $t = a + b$. The result is trivial for $t = 1$.

Given $t = a + b > 1$, we generate distribution $M[a, b]$ by constructing a value α , then with probability α we choose from the distribution of sequences (i.e. specifying an order of allocations) $M[a - 1, b]$ and append A , or else with probability $1 - \alpha$ we choose from the distribution $M[a, b - 1]$ and append B . We must show the existence of some α value such that the three conditions required by Lemma 1 will hold. \square

Conditions 2 and 3 of the lemma describe an interval in which the value $w^A(a, b)$ must fall, call it $I_m^{a,b}$. That is,

$$I_m^{a,b} = [w^A(a - 1, b), w^A(a, b - 1) + \Delta^{\oplus B}(a, b)].$$

Claim 1 shows that this interval is non-empty.

Claim 1 $w^A(a - 1, b) \leq w^A(a, b - 1) + \Delta^{\oplus B}(a, b)$.

Proof This follows by induction using condition 1 of the Lemma, which implies $w^A(a - 1, b) \leq w^A(a, b - 1) \leq w^A(a, b - 1) + \Delta^{\oplus B}(a, b)$. \square

Let $W_1^A(a, b)$ (respectively, $W_1^B(a, b)$) denote the expected payoff of agent A (respectively, agent B) if we let $\alpha = 1$. That is, $W_1^A(a, b)$ is the expected influence of agent A if we select a permutation from $M[a - 1, b]$ and append A , then use this permutation when applying our greedy algorithm. We define $W_0^A(a, b)$ and $W_0^B(a, b)$ similarly for $\alpha = 0$. The following claim follows from the adverse competition assumption.

Claim 2 $W_1^A(a, b) \geq w^A(a - 1, b)$ and $W_0^A(a, b) \leq w^A(a, b - 1)$.

Proof The first part of the claim follows because, for each fixed ordering in the support of $M[a - 1, b]$, appending an A to that ordering can only increase the welfare of agent A . Likewise, the second part of the claim follows because, for each ordering in the support of $M[a, b - 1]$, appending a B can only decrease the welfare of agent A (due to adverse competition). \square

We think of $W_1^A(a, b)$ and $W_0^A(a, b)$ as the influence for agent A for distributions that we can construct. Let $I_c^{a,b}$ denote the interval between $W_1^A(a, b)$ and $W_0^A(a, b)$. Note that we do not know which of $W_1^A(a, b)$ or $W_0^A(a, b)$ is greater. Claim 2 implies that:

Claim 3 $I_m^{a,b} \cap I_c^{a,b} \neq \emptyset$.

Proof It cannot be that $I_c^{a,b}$ lies entirely above $I_m^{a,b}$, since $W_0^A(a, b) \leq w^A(a, b - 1) \leq w^A(a, b - 1) + \Delta^{\oplus B}(a, b)$. Also, it cannot be that $I_c^{a,b}$ lies entirely below $I_m^{a,b}$, since $W_1^A(a, b) \geq w^A(a - 1, b)$. Thus $I_m^{a,b} \cap I_c^{a,b} \neq \emptyset$. \square

We can therefore write $I^{a,b} = I_m^{a,b} \cap I_c^{a,b}$. Note that any point in $I^{a,b}$ corresponds to a lottery (choosing any of the interval’s endpoint with some probability) we can construct for $M[a, b]$, which will satisfy conditions 2 and 3 of our Lemma. It remains

to show that we can choose this point so that condition 1 of Lemma 1 will also be satisfied. Our claim is that if we always choose α so that $w^A(a, b)$ is the minimum endpoint of $I^{a,b}$, then condition 1 will be satisfied.

With the above in mind, we will set

$$\alpha = \arg \min_{\alpha \in [0,1]} \left\{ \alpha W_1^A(a, b) + (1 - \alpha) W_0^A(a, b) \in I^{a,b} \right\} \tag{3}$$

Note that if we use this value of α to randomize between appending A to a permutation drawn from $M[a - 1, b]$ and appending B to a permutation from $M[a, b - 1]$, then the resulting value of $w^A(a, b)$ will indeed be $\min I^{a,b}$.

For all $a' + b' = t$, define $M[a', b']$ as described above. We now argue that this choice satisfies condition 1 of Lemma 1.

Claim 4 *If $a \geq 1$ then $w^A(a, b) \geq w^A(a - 1, b + 1)$.*

Proof Note first that $w^A(a, b) \geq w^A(a - 1, b)$, since $w^A(a, b) \in I_m^{a,b}$. Consider now the value of $w^A(a - 1, b + 1)$, which is the minimum of $I_c^{a-1,b+1} \cap I_m^{a-1,b+1}$. We will now bound the value of $w^A(a - 1, b + 1)$, by providing an upper bound on both the minimal endpoint of $I_c^{a-1,b+1}$ and the minimal endpoint of $I_m^{a-1,b+1}$.

For budgets $(a - 1, b + 1)$, the lower endpoint of $I_m^{a-1,b+1}$ is $w^A(a - 2, b + 1)$. On the other hand, $I_c^{a-1,b+1}$ contains point $W_0^A(a - 1, b + 1)$, which is the influence to agent A when we choose a permutation according to $M[a - 1, b]$ and append a ‘B’. However, since allocating an additional item to agent B in any fixed allocation can only degrade agent A ’s payoff, it must be that $W_0^A(a - 1, b + 1) \leq w^A(a - 1, b)$.

Thus the lower endpoint of $I_m^{a-1,b+1} \cap I_c^{a-1,b+1}$ is at most $\max\{w^A(a - 2, b + 1), w^A(a - 1, b)\}$. But $w^A(a - 2, b + 1) \leq w^A(a - 1, b)$ by induction (using condition 1 of Lemma 1).

We therefore conclude $w^A(a - 1, b + 1) \leq \max\{w^A(a - 2, b + 1), w^A(a - 1, b)\} \leq w^A(a - 1, b) \leq w^A(a, b)$, as required. \square

We have shown that table M can be filled with distributions that satisfy the conditions of Lemma 1. It remains to discuss the complexity of computing the entries of M . To this point we have not bounded the size of our distributions’ supports. We will modify the argument to show that the number of permutations required for each table entry $M[a, b]$ can be limited to only three, by induction on t .

Consider the distribution constructed for $M[a, b]$. By the inductive hypothesis, since $M[a - 1, b]$ and $M[a, b - 1]$ each have support of size 3, the support of this distribution has size at most 6: the three permutations in the support of $M[a - 1, b]$ with A appended, plus the three permutations in the support of $M[a, b - 1]$ with B appended. Each of these six permutations determines an allocation, say $(S_1, T_1), \dots, (S_6, T_6)$. For each allocation, we consider the two-dimensional point $(f_A(S_i, T_i), f_B(S_i, T_i))$ representing the welfare to A and B for the given allocation. We can interpret our construction of $M[a, b]$ as implementing a point $(w^A(a, b), w^B(a, b))$ with certain properties, such that this point lies in the convex hull of the six points $(f_A(S_1, T_1), f_B(S_1, T_1)), \dots, (f_A(S_6, T_6), f_B(S_6, T_6))$.

We now use the following well-known theorem [24]:

Lemma 2 (Carathéodory’s theorem) *Given a set $V \subset \mathbb{R}^n$ and a point $p \in \text{Conv}V$ in the convex hull of V , there exists a subset $A \subset V$ such that $|A| \leq n + 1$ and $p \in \text{Conv}A$.*

It must therefore be that our point $(w^A(a, b), w^B(a, b))$ lies in the convex hull of at most three of the points $(f_A(S_1, T_1), f_B(S_1, T_1)), \dots, (f_A(S_6, T_6), f_B(S_6, T_6))$. It follows that there exists a distribution with a support that consists of three of the six permutations corresponding to (a, b) . Finding this distribution can be done in constant time by considering $\binom{6}{3}$ sets of three allocations.⁸ We can therefore construct $M[a, b]$ as a distribution over at most 3 permutations, concluding the proof of Lemma 1.

The proof of Lemma 1 is constructive: it implies a recursive method for constructing the table M of distributions. That is, the procedure *ConstructDistributions* from Algorithm 1 (with input (a, b)) will proceed by filling table M in increasing order of t , up to $a + b$, by choosing the value of α for each table entry as in the proof of Lemma 1, then storing the implied distribution over three permutations. Note that we can explicitly store the allocations corresponding to the permutations in the table, making it simple to compute the submodular function values needed to determine α (which is stored as well). We conclude, given this implementation of *ConstructDistributions*, that Algorithm 1 is a polytime strategyproof 2-approximation to the 2-agent influence maximization problem.

4 A Strategyproof Mechanism for Three or More Players

To construct a strategyproof mechanism for $k > 2$ players, beyond the MeI and adverse competition assumptions, we will impose the additional agent indifference (AgI) property of the influence functions f_1, \dots, f_k . We show that, under these assumptions, there is a natural mechanism that is strategyproof when there are at least three players. In fact, it turns out that having three or more players in such a setting allows for a much simpler ordering than in the mechanism for the case of only two players⁹.

We note that the models for competitive influence spread proposed by Carnes et al. [10] and Bharathi et al. [11] are based on a cascade model of influence spread, and satisfy both the MeI and AgI assumptions.

⁸ Note that all quantities in this geometric problem are rational numbers, which are constructed via the sequence of operations described in the proof above and therefore have polynomial bit complexity. We can therefore solve the convex hull tasks described in this operation in polynomial time.

⁹ At this point, the reader may wonder if the two player case can be reduced to the case $k > 2$ by adding dummy agents with budget 0. This does not work because strategyproofness is defined over the space of *all possible* agent bids, so we cannot restrict our attention only to profiles in which some players bid 0. Our examples in Appendix “Counterexamples When There are Two Agents (Extended Discussion)” show that this is not just a nuance of the proof but rather an intrinsic obstacle to using the uniform distribution.

4.1 The Uniform Random Greedy Mechanism

Consider Algorithm 2, which we refer to as the uniform random greedy mechanism. This mechanism proceeds by first greedily selecting which subset of the ground set elements to allocate. It then chooses an ordering of the players’ bids uniformly at random from the set of all possible orderings, then assigns the selected elements to the players in this randomly chosen order. Note that this always results in a disjoint allocation.

Algorithm 2: Uniform Random Greedy Mechanism

```

Input: Ground set  $U = \{e_1, \dots, e_m\}$ , budget profile  $\mathbf{b}$ 
Output: An allocation profile  $\mathbf{S}$ 

1 Initialize:  $S_i \leftarrow \emptyset, i \leftarrow 0, j \leftarrow 0, I \leftarrow \emptyset, t \leftarrow \sum_i b_i$ ;
   /*Choose elements to assign. */
2 while  $i < t$  do

3    $u_i \leftarrow \arg \max_{c \in U} \{f(I \cup \{c\}) - f(I)\}$ ;
4    $I \leftarrow I \cup \{u_i\}; i \leftarrow i + 1$ ;
5 end
   /*Partition elements of  $I$ . */
6  $\Gamma \leftarrow \{\beta : [t] \rightarrow [k] \text{ s.t. } |\beta^{-1}(i)| = b_i \text{ for all } i\}$ ;
7 Choose  $\beta \in \Gamma$  uniformly at random;
8 while  $j < t$  do
9    $S_{\beta(j)} \leftarrow S_{\beta(j)} \cup \{u_j\}$ ;
10   $j \leftarrow j + 1$ ;
11 end
    
```

The MeI assumption implies that the random greedy mechanism obtains a constant factor approximation to the optimal social welfare. We now claim that, under the MeI and AgI assumptions, this mechanism is strategyproof as long as there are at least 3 players.

Theorem 2 *If there are $k \geq 3$ players and the AgI and MeI assumptions hold, then Algorithm 2 is a strategyproof mechanism. Furthermore, Algorithm 2 approximates the social welfare to within a factor of $\frac{e}{e-1}$ from the optimum.*

Proof As before, notice that lines 2–5 are an implementation of the standard greedy algorithm for maximizing a non-decreasing, submodular set-function subject to a uniform matroid constraint, as described in [17,23], and hence gives the specified approximation ratio.

Next, we show that Algorithm 2 is strategyproof. Fix bid profile \mathbf{b} and let $t = \sum_i b_i$. Let I be the union of all allocations made by Algorithm 2 on bid profile \mathbf{b} ; note that I depends only on t , due to MeI. Furthermore, each agent i will be allocated a uniformly random subset of I of size b_i . Thus, the expected utility of agent i can be expressed as a function of b_i and t , due to AgI. We can therefore write $w^i(b, t)$ for the expected utility of agent i when $b_i = b$ and $\sum_j b_j = t$ (recall that we let $w(t)$ denote the total social welfare when $\sum_i b_i = t$).

We now claim that $w^i(b, t) = \frac{b}{t}w(t)$ for all i and all $0 \leq b \leq t$. Note that this implies the desired result, since if our claim is true then for all i and all $0 \leq b \leq t$ we will have

$$w^i(b, t) = \frac{b}{t}w(t) \leq \frac{b+1}{t+1}w(t+1) = w^i(b+1, t+1)$$

which implies the required monotonicity condition.

It now remains to prove the claim. The adverse competition assumption implies that $w^i(0, t) \leq w^i(0, 0) = 0$ for all i and t . We next show that $w^i(1, t) = w^j(1, t)$ for all i, j , and $t \geq 1$. If $t = 1$ then this follows from the MeI assumption. So take $t \geq 2$ and pick any three agents i, j , and ℓ . Then, by the AgI assumption, we have

$$w^i(1, t) = w(t) - w^\ell(t-1, t) = w^j(1, t).$$

We next show that $w^i(b, t) = w^i(1, t) + w^i(b-1, t)$ for all i , all $b \geq 2$, and all $t \geq b$. Pick any three agents i, j , and ℓ , any $b \geq 2$, and any $t \geq b$. Then, by the AgI assumption,

$$\begin{aligned} w^i(b, t) &= w(t) - w^\ell(t-b, t) \\ &= w(t) - [w(t) - w^i(b-1, t) - w^j(1, t)] \\ &= w^i(b-1, t) + w^j(1, t) \\ &= w^i(b-1, t) + w^i(1, t). \end{aligned}$$

It then follows by simple induction that $w^i(b, t) = bw^i(1, t)$ for all $1 \leq b \leq t$. But now note that $w(t) = w^i(1, t) + w^j(t-1, t) = tw^i(1, t)$, and hence $w^i(1, t) = \frac{1}{t}w(t)$ and therefore $w^i(b, t) = \frac{b}{t}w(t)$ for all $0 \leq b \leq t$, as required. \square

Note that the proof of Theorem 2 makes crucial use of the fact that there are at least three players. Indeed, in Appendix “Counterexamples When There are Two Agents (Extended Discussion)” we give an example satisfying the MeI and (vacuously) AgI assumptions for which the random greedy algorithm is not strategyproof for two players.

5 Conclusions

We have presented a general framework for mechanisms that allocate items given an underlying submodular process. Although we have explicitly referred to spread processes over social networks, we only require oracle access to the outcome values, and thus our methods apply to any similar settings which uphold the properties we have required from the processes. We build on natural greedy algorithms to construct efficient strategyproof mechanisms that guarantee constant approximations to the social welfare.

The obvious open questions concern the removal of the MeI and AgI assumptions. For the case of $k = 2$ agents, is the MeI assumption necessary? For $k > 2$ agents, can

we eliminate either or both of the MeI and AgI assumptions? As discussed in Appendix “Tightness of Approach: More than Two Players”, it seems that a fundamentally new approach would be required to obtain an $O(1)$ -approximate strategyproof mechanism for $k > 2$ players.

Another natural and challenging extension would be to assume that nodes have costs for being initially allocated and then replace the cardinality constraint on each agent by a knapsack constraint. To do so, the most direct approach would be to try to utilize the known approximation for maximizing a non decreasing submodular function subject to one [25] or multiple [26] knapsack constraints. These methods do not seem to readily lend themselves to the approach we have been able to exploit in the case of cardinality constraints. We have also assumed a “demand satisfaction” condition dictating that each agent must receive their reported budget for initial adopters. Demand satisfaction is appropriate in settings where the mechanism can be seen to be fulfilling contracts. Can we relax this condition in some meaningful way? Without the demand satisfaction condition, it is trivial to achieve a strategyproof k approximation by allocating all initial elements to the agent who can achieve the most utility. Can we improve upon this trivial approximation factor without imposing the MeI and AgI assumptions?

Acknowledgements We would like to thank Yuval Filmus and anonymous referees for many helpful comments and discussions. In particular, we are indebted to Eyal Shani and one of the referees for pointing out a flaw in a previous version of this paper as regards the necessity of the MeI assumption in Theorem 1.

Appendix 1: Relation with Other Diffusion Models

In our results, we have made a number of modelling assumptions about agent utilities and social welfare. To some extent, we can argue that these assumptions may be necessary to be able to obtain truthfulness and constant approximation on the social welfare. Furthermore, we now provide some background on the relevance of our assumptions to the existing work on influence diffusion in social networks, which served as the running example throughout the paper.

Non-decreasing and Submodular Utilities and Social Welfare

To the best of our knowledge, in order to establish a constant approximation on the social welfare, all of the known models in competitive and non-competitive diffusion assume that the overall expected spread is a non-decreasing and submodular function with respect to the set of initial adopters. A main part of the seminal work by Kempe et al. [6] is the proof that the expected spread of two models of non-competitive diffusion process is indeed non-decreasing and submodular. This was later extended to more general processes in [7]. In the case of the competitive influence spread models in [10–12], it is shown that a player’s expected spread is a non-decreasing and submodular function of his initial set of nodes, while fixing the competitors’ allocations of nodes. This also implies that the total influence spread is a non-decreasing and submodular. Without any assumption on the nature of the social welfare function, it is NP-hard to obtain any non trivial approximation on the social welfare even for a single player.

Adverse Competition

In the initial adoption of (say) a technology, a competitor can indirectly benefit from competition so as to insure widespread adoption of the technology. However, once a technology is established (e.g., cell phone usage), the issue of influence spread amongst competitors should satisfy adverse competition. The same can be said for selecting a candidate in a political election. We also note that the previous competitive spread models [10–12] mentioned above also satisfy adverse competition. In its generality, the Goyal and Kearns model need not satisfy this assumption, but in order to obtain their positive result on the price of anarchy, they adopt a similar restriction (namely, that the adoption function at every node satisfies the condition that a player's probability of influencing an adjacent node cannot decrease in the absence of other players competing).

Furthermore, a simple example shows that the assumption of adverse competition is necessary for truthfulness. Consider the following two-player setting. The ground set is composed of two items: u_1 , which contributes a value of 1 to the receiving player and a value of N to her competitor (who did not receive u_1), and item u_2 which gives both players a value of 1. Now, consider the outcome of any mechanism when the bid profile is $(1, 1)$. Without loss of generality, one player, say player A , will receive u_1 , while the other player will get u_2 . The valuations would therefore be 2 and $N + 1$ for players A and B , respectively. In that case, player A would prefer to lower her bid to 0, which would guarantee her a valuation of N (player B would have to get u_1 , as otherwise the approximation ratio of the social welfare is unbounded as N grows). We conclude that unless the competition assumption holds, no strategyproof mechanism can, in general, obtain a bounded approximation ratio to the optimal social welfare. Although the example refers to deterministic allocations, the same argument can be made for randomized allocations.

Mechanism and Agent Indifference

In both the Wave Propagation model and the Distance-Based model presented in [10], the propagation of influence upholds both the mechanism and agent indifference properties. Goyal and Kearns [13] assume that the switching function that determines the probability that a node will adopt some technology is a function of the fraction of influenced neighbours (regardless of their assumed technologies). This immediately implies MeI, as these switching functions do not depend on the technologies involved. However, we note that whether or not the model satisfies the AgI assumption depends on the manner by which the nodes select specific technologies. The functions that determine this are termed selection functions in [13]. For example, suppose α_j is the fraction of nodes infected by technology j . If the switching function determining the probability of adopting technology i is $s(i) = \frac{\alpha_i}{\sum_j \alpha_j}$, then the model satisfies AgI. On

the other hand, if $s(i) = \frac{\alpha_i^2}{\sum_j \alpha_j^2}$, then the model would not satisfy AgI when there are more than 2 technologies.

Multiset Allocations and Disjointness

Our model assumes that each agent can be allocated a node at most once, and indeed most influence models assume that a node is allocated at most once in any initial allocation. However, we can extend our model to allow a node to be allocated multiple times to the same agent, as in (for example) the model of Goyal and Kearns [13]. To implement such an extension, we can simply consider a modified problem instance with many identical copies of a given node, treating each copy as a distinct element, and then proceed as though each element can be allocated at most once. The output of Algorithm 1 for two players and Algorithm 2 for more than two players would then be a profile of multisets with regard to the original network model. We note, however, that for the case of Algorithm 2, the MeI and AgI definitions effectively imply that if multisets are permitted, then non-disjoint allocations must be permitted as well, as the conditions cannot distinguish between an element being allocated to one agent twice or to two different agents.

Generality of the Model

A few words are in order about the generality of the model of diffusion under which we prove that Algorithm 1 is strategyproof and provides a $\frac{e}{e-1}$ -approximation. As noted, with the exception of the OR model, the analysis in previous competitive influence models assumes anonymous agents. Our general model does not require anonymity and hence we can accommodate agent specific edge weights (e.g. in determining the probability that influence is spread across an edge, or for determining whether the weighted sum of influenced neighbors crosses a given threshold of adoption). Our model also notably allows agent-independent node weights, for determining the value of an influenced node. Moreover, our abstract model does not specify any particular influence spread process, so long as the social welfare function is monotone submodular and each player's payoff is monotonically non-decreasing in his own set and non-increasing in the allocations to other players. In particular, our framework can be used to model probabilistic cascades as well as submodular threshold models.

Appendix 2: Counterexamples When There are Two Agents (Extended Discussion)

The locally greedy algorithm is defined over an *arbitrary* permutation of the allocation turns. At the core of our work, we seek to carefully construct such orderings in a manner that induces strategyproofness. We demonstrate that this algorithm due to Nemhauser et al [9] (see also Goundan and Schultz [17]) is not, in general, strategyproof for some natural methods for choosing the ordering of the allocation between two agents.

To clarify the context when there are only two agents, we refer to them as agent A and agent B and their utilities as f_A and f_B respectively. We give examples of a set U and functions f_A and f_B (satisfying the conditions of our model) such that natural greedy algorithms for choosing sets S and T result in non-monotonicities. Our

examples will all easily extend to the case of $k > 2$ agents (but not satisfying agent indifference).

The OR Model

We will consider examples of a special case of the OR model for influence spread, as studied in [12]. Let $G = (V, E)$ be a graph with fractional edge-weights $p : E \rightarrow [0, 1]$, vertex weights w_v for each $v \in V$, and sets $I_A, I_B \subseteq V$ of “initial adopters” allocated to each player. We use vertex weights for clarity in our examples; in Appendix “Counterexamples with Unweighted Nodes” we show how to modify the examples given in this section to be unweighted. We emphasize that all “infected” nodes (including any initially selected) contribute their weight to the expected social welfare and individual values of the players. The process unfolds in discrete steps. For each $u_A \in I_A$ and v_A such that $(u_A, v_A) \in E$, u_A , once infected, will have a single chance to “infect” v_A with probability $p(u_A, v_A)$. Define the same, single-step process for the nodes in I_B , and let O_A and O_B be the nodes infected by nodes in I_A and I_B , respectively. Note that the infection process defined for each individual player is an instance of the Independent Cascade model as studied by Kempe et al. [6]. Finally, nodes that are contained in $O_A \setminus O_B$ will be assigned to player A , nodes in $O_B \setminus O_A$ will be assigned to B , and any nodes in $O_A \cap O_B$ will be assigned to one player or the other by flipping a fair coin.

In our examples, we consider two identical players each having utility equal to the weight of the final set of nodes assigned by the spread process. It can be easily verified that both the expected social welfare (total weight of influenced nodes) and the expected individual values (fixing the other player’s allocation) are submodular set-functions.

Deterministic Greedy Algorithms that are Not Strategyproof

We demonstrate that the more obvious deterministic orderings for the greedy algorithm fail. First, consider the “dictatorship” ordering, in which (without loss of generality by symmetry) player A is first allocated nodes according to his budget, and only then player B is allocated nodes. Our example showing non-truthfulness also applies to an ordering that would always select the player having the largest current unsatisfied budget breaking ties (again without loss of generality by symmetry) in favor of player A . Consider the graph depicted in Fig. 2a. When player A bids 1 and player B bids 1 as well, the algorithm will allocate c_1 to player A , as it contributes the maximal marginal gain to the social welfare, and will allocate c_3 to player B . The value of the allocation for player A is 2.

However, notice that if player A increases its bid to 2, the mechanism will allocate nodes c_1 and c_3 to player A , and allocate c_2 to B . In this case player A receives an extra value of $\frac{1}{2}$ from node c_3 , but the allocation of c_2 to B will “pollute” player A ’s value from c_1 : he will receive nodes u_1 and u_2 each with probability $\frac{1}{10} + \frac{1}{2} \cdot \frac{9}{10} = \frac{11}{20}$.




Figure 2 Counter-examples for the mechanism under the deterministic dictatorship and Round Robin orderings. In both cases, we set the weights $w_{c_i} = 0$ and $w_{u_i} = 1$, for all $1 \leq i \leq 4$. Additionally, we let $0 < \epsilon < \frac{1}{8}$. **a** The counter-example for the deterministic mechanism with a dictatorship ordering. The initial budget for both players is 1. **b** The counter-example for the deterministic algorithm under a Round Robin ordering. The initial budgets for players A and B are 1 and 2, respectively

Thus the total expected value for player A is only $\frac{16}{10}$, and hence the algorithm is non-monotone in the bid of player A.

Next, consider the Round Robin ordering, in which the mechanism alternates between allocating a node to player A and to player B. Our example here also applies to the case when the mechanism always chooses the player having the smallest current unsatisfied budget breaking ties in favor of player A. Consider the instance given in Fig. 2b. When the bids of players A and B are 1 and 2, respectively, the algorithm will first allocate c_1 to player A, and then it will subsequently allocate nodes c_3 and c_4 to player B, which results in a payoff of 1 for player A. If player A were to increase his bid to 2, then the mechanism would allocate nodes c_1 and c_4 to player A, and nodes c_2 and c_3 to player B, for a payoff of $3 \cdot \epsilon + 2 \cdot \epsilon + (1 - 2 \cdot \epsilon) \cdot \frac{1}{2} = \frac{1}{2} + 4 \cdot \epsilon < 1$ (since $0 < \epsilon < \frac{1}{8}$). Therefore, the monotonicity is violated for the payoff to player A.


The Uniform Random Greedy Algorithm is Not Strategyproof

As we shall see in Sect. 4, for the case of $k > 2$ agents in the restricted setting that assumes mechanism and agent indifference, a very simple mechanism admits a strategyproof mechanism that provides an $\frac{e}{e-1}$ approximation to the optimal social welfare. More specifically, we show that under these assumptions on the social welfare agent utilities, taking a uniformly random permutation over the allocation turns is a strategyproof algorithm. In contrast, for the case of $k = 2$, and even with these additional restrictions (although the agent indifference assumption turns out to be vacuous in this case), the uniformly random mechanism is not strategyproof.

Consider the example given in Fig. 3. We note that for this example, Algorithm 2 in Sect. 4 is equivalent to first choosing a random order of allocation (e.g. choosing all possible permutations satisfying agent budgets with equal probability) and then allocating greedily. The greedy algorithm will allocate one of c_2, c_3, c_4 and c_5 to one of the players, then allocate c_1 , and then any remaining nodes.

Let player A’s budget be 3 and player B’s budget be 1. In this case, with probability $\frac{1}{4}$, player B will be allocated c_1 (i.e. when B’s allocation occurs second), in which case player A’s expected value would be 1. Also, with probability $\frac{3}{4}$, player B will be

Figure 3 The counterexample for the mechanism that allocated according to a random ordering of the turns ($0 < \epsilon \ll 1$).
 $w_{c_i} = \epsilon, i = 1, \dots, 5,$
 $w_{u_i} = 1, i = 1, 2$



allocated one of $\{c_2, c_3, c_4, c_5\}$, in which case player A 's expected outcome would be $\frac{1}{2} + \epsilon$. In total, player A 's expected payoff will be $\frac{5}{8} + \frac{3}{4}\epsilon$.

If player A were to increase his budget to 4, then with probability $\frac{1}{5}$ player B will be allocated c_1 , in which case player A 's outcome will be 1. On the other hand, player A 's expected payoff will be $\frac{1}{2} + \epsilon$ if B is allocated one of $\{c_2, c_3, c_4, c_5\}$, which occurs with probability $\frac{4}{5}$. In total, player A 's expected outcome will be $\frac{3}{5} + \frac{4}{5}\epsilon < \frac{5}{8} + \frac{3}{4}\epsilon$, implying that this algorithm is non-monotone.

Appendix 3: Counterexamples with Unweighted Nodes

In Sect. “Counterexamples When There are Two Agents (Extended Discussion)” we constructed specific examples of influence spread instances for the OR model, to illustrate that simple greedy methods are not necessarily strategyproof for the case of two players. These examples used weighted nodes which our model allows. For the sake of completeness, we now show that these examples can be extended to the case of unweighted nodes.

We focus on the example from Sect. “The Uniform Random Greedy Algorithm is Not Strategyproof” to illustrate the idea; the other examples can be extended in a similar fashion. In that example there were nodes u_1 and u_2 of weight 1, and nodes c_1, \dots, c_5 of weight 0. We modify the example as follows. We choose a sufficiently large integer $N > 1$ and a sufficiently small $\epsilon > 0$. We will replace node u_1 with a set S of N independent nodes. We replace the ϵ -weighted edge from c_1 to u_1 with an ϵ -weighted edge from c_1 to each node in T .

Similarly, we replace u_2 by a set T of N independent nodes. For each c_i , we replace the unit-weight edge from c_i to u_2 with a unit weight edge from c_i to each node in T .

In this example, if the sum of agent budgets is at most 5, the greedy algorithm will never allocate any nodes in S or T . The allocation and analysis then proceeds just as in Sect. “The Uniform Random Greedy Algorithm is Not Strategyproof”, to demonstrate that if agent B declares 1 then agent A would rather declare 3 than 4.

Appendix 4: Tightness of Approach: More than Two Players

The mechanism we construct in Sect. 3.2 is applicable to settings in which there are precisely two competing players, and our mechanism in Sect. 4 for more than three players requires the MeI and AgI assumptions. A natural open question is whether these results can be extended to the general case of three or more agents without the

MeI and AgI restrictions. In this section we briefly describe the difficulty in applying our approach to settings with three players.

For the case of two players in Sect. 3.2, our mechanism was built from an initial greedy algorithm by randomizing over orderings under which to assign elements to players. Our construction is recursive: we demonstrated that if we can define the behaviour of a strategyproof mechanism for all possible budget declarations up to a total of at most t , then we can extend this to a strategyproof mechanism for all possible budget declarations that total at most $t + 1$. A key observation that makes this extension possible is the direct relation between the utilities of the two players. This manifests itself in the cross monotonicities that we utilize in the inductive argument. In addition, the strategyproofness condition (i.e. agent monotonicity) can be equivalently re-expressed as a certain “budget competition” property: if one player increases his budget, then the expected utility for the other player cannot increase by more than the marginal gain the total welfare. In other words, for all $a + b \geq 1$, a strategyproof mechanism must satisfy $w^A(a, b) - w^A(a, b - 1) \leq \Delta^{\oplus B}(a, b)$ where $\Delta^{\oplus B}(a, b) = w(a, b) - w(a, b - 1)$ and a similar consequence with regard to $\Delta^{\oplus A}(a, b)$.

Claim 5 (Equivalence of monotonicity and budget competition for two players)

$$1: w^A(a, b) - w^A(a, b - 1) \leq w(a, b) - w(a, b - 1) \text{ iff } w^B(a, b - 1) \leq w^B(a, b)$$

$$2: w^B(a, b) - w^B(a - 1, b) \leq w(a, b) - w(a - 1, b) \text{ iff } w^A(a - 1, b) \leq w^A(a, b)$$

Proof We provide the proof for $\Delta^{\oplus B}(a, b)$.

$$w^B(a, b - 1) \leq w^B(a, b) \text{ iff}$$

$$w^A(a, b) + w^B(a, b - 1) \leq w^A(a, b) + w^B(a, b) \text{ iff}$$

$$w^A(a, b) + w^B(a, b - 1) + w^A(a, b - 1) \leq w^A(a, b) + w^B(a, b) + w^A(a, b - 1) \text{ iff}$$

$$w^A(a, b) + w(a, b - 1) \leq w(a, b) + w^A(a, b - 1) \text{ iff}$$

$$w^A(a, b) - w^A(a, b - 1) \leq w(a, b) - w(a, b - 1)$$

□

A direct extension of our approach to three players would involve proving inductively that an allocation rule that satisfies these conditions for all budgets that total at most t can always be extended to handle budgets that total up to $t + 1$. We now give an example to show that this is not the case, even when our underlying submodular function takes a very simple linear form.

Suppose we have three players A, B , and C , and suppose our ground set U contains a single element g of value; all other elements are worth nothing. The utility for each agent is 1 if their allocation contains g , otherwise their utility is 0. In this case, the locally greedy algorithm simply gives element g to the first player that is chosen for allocation; the remaining allocations have no effect on the utility of any player. Note then that the marginal gain in social welfare is 1 for the first allocation, and 0 for all subsequent allocations made by the greedy algorithm.

We now define the behaviour of a mechanism for all budget declarations totalling at most 2. Note that the relevant feature of this mechanism is the (possibly randomized) choice of which agent is first in the order presented to the greedy algorithm. We present this behaviour in the following table.

Budgets (a, b, c)	Player selected	Utilities (w^A, w^B, w^C)
$(0, 0, 0)$	N/A	$(0, 0, 0)$
$(1, 0, 0)$	A	$(1, 0, 0)$
$(0, 1, 0)$	B	$(0, 1, 0)$
$(0, 0, 1)$	C	$(0, 0, 1)$
$(1, 1, 0)$	A	$(1, 0, 0)$
$(0, 1, 1)$	B	$(0, 1, 0)$
$(1, 0, 1)$	C	$(0, 0, 1)$

We note that this mechanism (restricted to these type profiles) is strategyproof, satisfies the budget competition property, and also satisfies the cross-monotonicity properties (i.e. in the invariants of Theorem 1). However, we claim that no allocation on input $(1, 1, 1)$ that obtains positive social welfare can maintain the budget competition property. To see this, note that the budget competition property would imply that $w^A(1, 1, 1) \leq w^A(1, 0, 1) + \Delta^{\oplus B}(1, 1, 1) = w^A(1, 0, 1) = 0$. Similarly, we must have $w^B(1, 1, 1) = w^C(1, 1, 1) = 0$. Thus, in order to maintain the budget competition property, our mechanism would have to generate social welfare 0 on input $(1, 1, 1)$, resulting in an unbounded approximation factor. We conclude that there is no way to extend this specific mechanism for budgets totalling at most 2 to a (strategyproof) mechanism for budgets totalling at most 3 while maintaining the constant approximation factor of the locally greedy algorithm.

Roughly speaking, the problem illustrated by this example is that the presence of more than two bidders means that a substantial increase in the utility gained by one player does not necessarily imply any limits in the utility of another specific player. This is in contrast to the case of two players, in which the utilities of the two players are more directly related. This fundamental difference seems to indicate that substantially different techniques will be required in order to construct strategyproof mechanisms with three or more players.

A different (and natural) approach would be to employ the solution for two players by grouping all but one player at a time, and running the mechanism for two players recursively. However, this method seems ineffective in our setting, as interdependencies between the players' outcomes can introduce non-monotonicities. This brings into question whether or not the locally greedy method can be made strategyproof by some method of randomizing over the order in which allocations are made.

This “2 vs 3 barrier” is, of course, not unique to our problem. Many optimization problems (such as graph coloring) are easily solvable when the size parameter is $k = 2$ but become NP-hard when $k \geq 3$. Closer to our setting, the 2 vs 3 barrier has been discussed in recent papers concerning mechanism design without payments, such as in the Lu et al. [27] results for k -facility location. Additionally Ashlagi et al. discussed similar issues [28] in the context of mechanisms for kidney exchange. They show

that for n points on the line, there is a deterministic (respectively, randomized) strategyproof mechanism for placing $k = 2$ facilities (so as to minimize the sum of distances to the nearest facility) with approximation ratio $n - 2$ (respectively, 4) whereas for $k = 3$ facilities, they do not know if there is any bounded ratio for deterministic strategyproof mechanisms and the best known approximation for randomized strategyproof mechanisms is $O(n)$.

References

1. Granovetter, M.: Threshold models of collective behavior. *Am. J. Soc.* **83**, 1420–1443 (1978)
2. Schelling, T.: *Micromotives and Macrobehavior*. Norton, New York (1978)
3. Brown, J.J., Reingen, P.H.: Social ties and word of mouth referral behavior. *J. Consum. Res.* **14**, 350–362 (1987)
4. Goldenberg, J., Libai, B., Mulle, E.: Talk of the network: a complex systems look at the underlying process of word-of-mouth. *Mark. Lett.* **12**, 211–223 (2001)
5. Centola, D., Macy, M.: Complex contagions and the weakness of long ties. *Am. J. Sociol.* **113**, 702–734 (2007)
6. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In: *KDD '03: Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 137–146. ACM, New York (2003)
7. Kempe, D., Kleinberg, J., Tardos, É.: Influential nodes in a diffusion model for social networks. In: *Caires, L., Italiano, G., Monteiro, L., Palamidessi, C., Yung, M. (eds.) Automata, Languages and Programming. Volume 4858 of Lecture Notes in Computer Science*, vol. 3580, pp. 1127–1138. Springer, Berlin (2005)
8. Mossel, E., Roch, S.: Submodularity of influence in social networks: from local to global. *SIAM J. Comput.* **39**, 2176–2188 (2010)
9. Nemhauser, G.L., Wolsey, L.A., Fisher, M.L.: An analysis of approximations for maximizing submodular set functions-II. *Math. Program.* **14**, 265–294 (1978)
10. Carnes, T., Nagarajan, C., Wild, S.M., van Zuylen, A.: Maximizing influence in a competitive social network: a follower's perspective. In: *Proceedings of the Ninth International Conference on Electronic Commerce. ICEC '07*, pp. 351–360. ACM, New York (2007)
11. Bharathi, S., Kempe, D., Salek, M.: Competitive influence maximization in social networks. In: *Deng, X., Graham, F. (eds.) Internet and Network Economics. Volume 4858 of Lecture Notes in Computer Science*, vol. 4858, pp. 306–311. Springer, Berlin (2007)
12. Borodin, A., Filmus, Y., Oren, J.: Threshold models for competitive influence in social networks. In: *Saber, A. (ed.) Internet and Network Economics. Volume 6484 of Lecture Notes in Computer Science*, vol. 6484, pp. 539–550. Springer, Berlin (2010)
13. Goyal, S., Kearns, M.: Competitive contagion in networks. In: *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing. STOC '12*, pp. 759–774. ACM, New York, NY, USA (2012)
14. Feldman, J., Mehta, A., Mirrokni, V., Muthukrishnan, S.: Online stochastic matching: beating $1-1/e$. *FOCS*, pp. 117–126 (2009)
15. Feldman, J., Henzinger, M., Korula, N., Mirrokni, V.S., Stein, C.: Online stochastic packing applied to display ad allocation. In: *Proceedings of the 18th Annual European Conference on Algorithms: Part I. ESA'10*, pp. 182–194. Springer, Berlin (2010)
16. Borodin, A., Braverman, M., Lucier, B., Oren, J.: Strategyproof mechanisms for competitive influence in networks. In: *WWW*, pp. 141–150 (2013)
17. Goundan, P., Schulz, A.: Revisiting the greedy approach to submodular set function maximization. *Optim. Online*, 1–25 (2007)
18. Dubey, P., Garg, R., De Meyer, B.: Competing for customers in a social network: the quasi-linear case. In: *Spirakis, P., Mavronicolas, M., Kontogiannis, S. (eds.) Internet and Network Economics. Volume 4286 of Lecture Notes in Computer Science*, vol. 4286, pp. 162–173. Springer, Berlin (2006)
19. Kostka, J., Oswald, Y.A., Wattenhofer, R.: Word of mouth: rumor dissemination in social networks. In: *Proceedings of the 15th International Colloquium on Structural Information and Communication Complexity. SIROCCO '08*, pp. 185–196. Springer, Berlin (2008)

20. Apt, K., Markakis, E.: Diffusion in social networks with competing products. In: Persiano, G. (ed.) *Algorithmic Game Theory*. Volume 6982 of *Lecture Notes in Computer Science*, vol. 6982, pp. 212–223. Springer, Berlin (2011)
21. Borodin, A., Braverman, M., Lucier, B., Oren, J.: Truthful mechanisms for competing submodular processes. *CoRR*. [arXiv:1202.2097](https://arxiv.org/abs/1202.2097) (2012)
22. Singer, Y.: How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In: *WSDM*, pp. 733–742 (2012)
23. Nemhauser, G.L., Wolsey, L.A., Fisher, M.L.: An analysis of approximations for maximizing submodular set functions-i. *Math. Program.* **14**, 265–294 (1978). doi:[10.1007/BF01588971](https://doi.org/10.1007/BF01588971)
24. Rockafellar, R.T.: *Convex Analysis* (Princeton Landmarks in Mathematics and Physics). Princeton University Press, Princeton (1996)
25. Sviridenko, M.: A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.* **32**, 41–43 (2004)
26. Kulik, A., Shachnai, H., Tamir, T.: Maximizing submodular set functions subject to multiple linear constraints. In: *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '09, pp. 545–554. Society for Industrial and Applied Mathematics, Philadelphia (2009)
27. Lu, P., Sun, X., Wang, Y., Zhu, Z.A.: Asymptotically optimal strategy-proof mechanisms for two-facility games. In: *Proceedings of the 11th ACM Conference on Electronic commerce*. EC '10, pp. 315–324. ACM, New York (2010)
28. Ashlagi, I., Fischer, F., Kash, I., Procaccia, A.D.: Mix and match. In: *Proceedings of the 11th ACM Conference on Electronic commerce*. EC '10, pp. 305–314. ACM, New York (2010)