

Strategyproof Mechanisms for Competitive Influence in Networks

Allan Borodin* Mark Braverman† Brendan Lucier‡ Joel Oren§

December 5, 2013

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 demands 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 not be incentivized to underreport their budget demands.

We show that when there are two players, the social welfare can be 2-approximated by a polynomial-time strategyproof mechanism. Our mechanism is defined recursively, randomizing the order in which advertisers are allocated seeds according to a particular greedy method. For three or more players, we demonstrate that under additional assumptions (satisfied by many existing models of influence spread) there exists a simpler strategyproof $\frac{e}{e-1}$ -approximation mechanism; notably, this second mechanism is not necessarily strategyproof when there are only two players.

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 [15, 25, 6, 12, 8], 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") [16, 17, 21]. One commonality among many of these (stochastic) models is that the expected number of eventual adopters is a non-decreasing submodular function of the seed set. This implies that natural greedy methods [22] can be used to choose initial adopters 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 who wishes to provide this service to its advertisers.

*Dept of Computer Science, University of Toronto, bor@cs.toronto.edu

†Dept of Computer Science, Princeton University, mbraverm@cs.princeton.edu

‡Microsoft Research, brlucier@microsoft.com

§Dept of Computer Science, University of Toronto, oren@cs.toronto.edu

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 demands (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 eccentricities 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 [7, 3, 5, 14]. 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 underreport their budget demands 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 [14, 3]. 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. In this sense our framework is closer in spirit to matching algorithms for display advertising [11, 10] 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 model of competitive influence spread is described formally in Section 2. Our formulation captures and extends many existing models of influence spread, allowing incorporation of features such as node weights, player-specific spread probabilities, and non-linear selection probabilities. A more detailed discussion appears in Appendix A.

We wish to design mechanisms that are strategyproof, in that rational agents are incentivized to truthfully reveal their demands. 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.

Our Results: We design three different strategyproof mechanisms for the competitive influence maximization problem, for use in varying circumstances. Our main result is a 2-approximate strategyproof mechanism for use when there are two competing advertisers, under a very general model of influence spread. 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 construction is based upon a greedy algorithm for submodular function maximization subject to a partition matroid constraint, known as the locally greedy algorithm [22, 13]. This algorithm repeatedly chooses an agent in each round, and assigns a node to that agent in order to maximize the marginal increase

to social welfare. As we discuss in Section 3, this algorithm is not strategyproof in general. However, it has the property that the choice of agent in each round is arbitrary; this provides a degree of freedom that can be exploited to obtain strategyproofness. Indeed, for the case of two agents, we show how to recursively construct a distribution over potential allocations returned by locally greedy algorithms, with the property that each agent’s expected individual value under this distribution is monotone¹ with respect to the number of initial elements allocated.

Our second mechanism is for three or more players, under some natural restrictions on the influence spread process. Specifically, we require two properties: first, the social welfare is independent of the manner in which elements are partitioned among the players (*mechanism indifference*). Second, 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*). These conditions are defined formally in Section 5. We note that these assumptions are implicit in many prior models of influence spread [7, 3]. Under these assumptions, we develop a strategyproof mechanism that obtains a $\frac{e}{e-1}$ -approximation to the optimal social welfare when there are three or more players. Interestingly, our analysis makes crucial use of the presence of three or more players, and indeed we show that this mechanism fails to be strategyproof when only two players are present, even with these two additional assumptions².

Our final mechanism construction satisfies an additional constraint that agent allocations be disjoint. In its most general form, our problem specification does not require that the set of elements allocated to the agents be disjoint³. Our first mechanism described above may place a given node in the seed sets of multiple players. Our second mechanism for more than two players produces a disjoint allocation when the greedy algorithm used for a single player results in a disjoint allocation. When it is desirable for allocations to be disjoint, we show how our construction can be modified to work under this additional requirement, resulting in a strategyproof 3-approximation mechanism. This result requires that we impose a symmetry assumption on the influence spread model, which states that the outcome of the influence process is invariant under relabelling of the players⁴.

Our mechanisms run in time polynomial in the demands submitted by the agents and in the size of the underlying ground set. This dependence on the demand 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 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 demands. Generally speaking, the spread process itself is randomized and as in [16, 17], the oracle can be viewed as providing an element that approximately maximizes the marginal gain by sampling enough trials of this process [16, 17]. Our analysis also holds when such approximate marginal maximizers are used to implement our underlying greedy algorithm; following the exposition in [13], such an approximate maximizer provides an approximation that approaches 2 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.

Related Work: Models of influence spread in networks, covering both cascade and threshold phenomena, are well-studied in the sociology and marketing literature [15, 25]. The (non-competitive) problem of maximizing influence in social networks was theoretically modelled by Kempe et al. [16, 17]. Subsequent papers extended these models to a competitive setting in which there are multiple advertisers. Carnes et al. [7] suggested the Wave Propagation model and the Distance Based model, which were based on the Independent Cascade model. Additionally, Dubey et al. [9], Bharathi et al. [3], Kosta et al. [18], and Apt et al. [1] 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

¹We use the word *monotone* in its game-theoretic sense, meaning that a player’s outcome is a monotone function of its bid. We distinguish this from the monotonicity of the social welfare function of the mechanism, and use the term *non-decreasing* when referring to the social welfare function.

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

³Many prior models of competitive influence do allow non-disjoint allocations [14, 3]; our intention is to demonstrate that a disjointness condition can be accommodated if necessary, rather than imply that non-disjointness is undesirable.

⁴We note that this property holds for most models of influence spread studied in the literature [14, 3, 7, 1].

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. [5] 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 work [4], Goyal and Kearns [14] 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 [21]), so our mechanism for two players applies to their model as well⁵. 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 [26] 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.

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 set⁶ S_i to each player i . For the most part we will follow the convention that these sets should be disjoint, though in general our model does not require disjointness. In particular, we consider a setting in which sets need not be disjoint in Section 4.

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 $S_i \subseteq S'_i$.

We impose one final model assumption, which we call *adverse competition*: that 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 of influence for another player. We discuss the motivation for and necessity of this assumption in Appendix A.

⁵An “adverse competition” assumption in [14] is stated for $k = 2$ agents and holds at every node. Their assumption is somewhat weaker than ours, which we only apply to the social welfare function. See section 2.

⁶For 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 A for further discussion.

We study the following algorithmic problem. Given input values $b_1, \dots, b_k \geq 0$, we wish to find sets $S_1, \dots, S_k \subseteq U$, with $|S_i| = b_i$, for all $i = 1, \dots, k$, such that $f(S_1, \dots, S_k)$ is maximized. We assume we are given oracle access to the functions f and f_1, \dots, f_k . Note that we impose a “demand satisfaction” condition on the mechanism, that each agent is allocated all of his demand. (To this end we will assume that $|U| \geq \sum_{i=1}^k b_i$; i.e. that there are enough items to allocate).

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 bid 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 player i has a *type* \tilde{b}_i , representing the maximum number of elements they can be allocated. The utility of player i for allocation $\mathbf{S} = (S_1, \dots, S_k)$ is

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

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.

The problem of maximizing welfare function $f(\cdot)$ subject to the reported demands can be stated in the framework of maximizing a submodular set-function subject to a *partition matroid* constraint. An instance of a partition matroid $\mathcal{M} = (E, \mathcal{F})$ is given by a union of disjoint sets $E = \bigcup_{i=1, \dots, k} E_i$, and a set of corresponding cardinality constraints d_1, \dots, d_k . A set X is in \mathcal{F} , i.e. is *independent*, if $|X \cap E_i| \leq d_i$, for all $1 \leq i \leq k$. That is, an independent set is formed by taking no more than the prescribed size constraint for each of the sets. The optimization problem to find an independent set that maximizes a non-decreasing and submodular set-function $g : \mathcal{F} \rightarrow \mathbb{R}_{\geq 0}$. Our problem falls into this framework by setting the ground set to be $U \times \{1, \dots, k\}$, the cardinality constraints $d_i = b_i$, for all i and setting the objective function to be the social welfare:

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

We note, however, that this formulation does not apply if the allocated sets are required to be disjoint. The addition of disjointness causes our constraint to no longer take the form of a matroid, an issue which will be addressed in Section 6. Also note that this alternative definition of our setting conforms to the single-parameter convention of submodular set-functions. However, we will mostly refer to the former formulation of the problem for clarity and succinctness.

As a result of this correspondence with the framework of partition matroids, we will be interested in a particular greedy algorithm for this algorithmic problem, known as a *locally greedy* algorithm, studied in [23], which was subsequently extended in [13]. The algorithm proceeds by fixing some arbitrary permutation of the multiset composed of b_i i 's for each player i . It then iteratively builds the allocation \mathbf{S} where, on iteration j , it chooses $u \in \arg \max_c \{f(S_i \cup \{c\}, \mathbf{S}_{-i}) - f(S_i, \mathbf{S}_{-i})\}$ and adds u to S_i , where i is the j th element of the permutation. Regardless of the permutation selected, this algorithm is guaranteed to obtain a 2-approximation to the optimal allocation subject to the given cardinality constraints [23, 13].

3 Counter examples when there are two agents

The locally greedy algorithm [22] (see also [13]) is defined over an *arbitrary* permutation of the agents allocation turns. In Section 4 we carefully construct such orderings in a manner that induces strategyproofness for two players. To motivate these algorithmic gymnastics, we now demonstrate that more natural orderings fail to result in strategyproofness.

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 process⁷ 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. 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 locally 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 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 orderings for the locally greedy algorithm fail to result in strategyproof mechanisms. Appendix B 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 Section 5 we show that for the case of $k > 2$ agents, in a setting that assumes two additional restrictions called MeI and AgI (which will be defined in Section 5), taking a uniformly random permutation over the allocation turns is a strategyproof algorithm and results in an $\frac{e}{e-1}$ approximation to the optimal social welfare. In contrast, for the case of $k = 2$, and even with these additional restrictions (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), the uniformly random mechanism is not strategyproof.

4 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 Section 5, under additional assumptions that are not necessary for the case $k = 2$. Our mechanism is based on the local greedy algorithm described in Section 2. We will focus on cases in which the allocations to the two agents need not be disjoint; in Section 6 we extend our result to handle disjointness constraints when agents are “anonymous.”

A nice property of the local greedy algorithm is that its worst-case approximation factor of 2 holds even if we arbitrarily fix the order in which allocations are made to players A and B . This grants a degree of freedom that we will use to satisfy strategyproofness. Given a particular pair of budgets (a, b) , we will randomize over

⁷This process is a simplification of the OR model [5].

possible orderings in which to allocate to the two agents, and then apply the greedy algorithm to whichever permutation we choose. The key to the algorithm will be the manner in which we choose the distribution to randomize over, which will depend on the declared budgets and the influence functions f_i . As it turns out, some of the more immediate ways of selecting an ordering lead to non-strategyproof mechanisms. See Appendix 3 for a survey of naïve orderings. Indeed, it is not even clear a priori that distributions exist that simultaneously monotize the expected allocation for both players. Our main technical contribution is a proof that such distributions do exist, and moreover can be explicitly constructed in polynomial time.

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 distributions for the case $t = 1$ (which are trivial), followed by $t = 2$, etc. To construct the distribution for demands (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 player B . If we choose the former 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 player A under the distribution for (a, b) is at least that of the distribution for $(a - 1, b)$, and similarly for player 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 each α . 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.

4.1 The Allocation Algorithm

Our algorithm will proceed by choosing a distribution over orders in which nodes are allocated to the two players. 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. That is, items are chosen for the players 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 specific algorithm to ensure strategyproofness. We note that this technique could be applied to *any* allocation algorithm with this property; we will make use of this observation in Section 6.

Recall that the approximation guarantee for the greedy allocation does not depend on the order of assignment implemented in lines 3-11, so that the allocation returned by the algorithm will be a 2-approximation to the optimal total influence regardless of the permutation chosen on line 2. It remains only to demonstrate that we can construct our distributions in such a way that the expected payoff to each player is monotone increasing in his bid.

Algorithm 1: Allocation Mechanism

Input: Ground set $U = \{e_1, \dots, e_n\}$, budgets a, b for players A and B , respectively
Output: An allocation $I_A, I_B \subseteq U$ for the two players

```
/* Build permutation table. */
1  $M \leftarrow \text{ConstructDistributions}(a, b)$ ;
/*  $M[a, b]$  will be a distribution over sequences  $(y_1, \dots, y_{a+b}) \in \{A, B\}^{a+b}$  */
2 Choose  $(y_1, \dots, y_{a+b})$  from distribution  $M[a, b]$ ;
3 for  $i = 1 \dots a + b$  do
4   if  $y_i = A$  then
5      $u \leftarrow \text{argmax}_{c \in U} \{f(I_A \cup \{c\}, I_B) - f(I_A, I_B)\}$ ;
6      $I_A \leftarrow I_A \cup \{u\}$ ;
7   else
8      $u \leftarrow \text{argmax}_{c \in U} \{f(I_A, I_B \cup \{c\}) - f(I_A, I_B)\}$ ;
9      $I_B \leftarrow I_B \cup \{u\}$ ;
10  end
11 end
```

4.2 Constructing matrix M

We describe the procedure *ConstructDistributions*, used in Algorithm 1, to generate distributions over orderings of assignments to players 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 α , $0 \leq \alpha \leq 1$, 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 $\Delta^{\oplus B}(a, b) = w(a, b) - w(a, b - 1)$. That is, $\Delta^{\oplus B}(a, b)$ is the marginal gain in total welfare when agent B increases his bid from $b - 1$ to b , given matrix M .

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 player B 's valuation is monotone increasing with his bid:

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

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.

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 2 shows that this interval is non-empty.

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

Proof. This follows by induction applied to 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 (respectively, W_1^B) denote the expected payoff of player A (respectively, player B) if we let $\alpha = 1$. That is, W_1^A is the expected influence of player 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 and W_0^B similarly for $\alpha = 0$. The following claim follows from the adverse competition assumption.

Claim 3. $W_1^A \geq w^A(a - 1, b)$ and $W_0^A \leq w^A(a, b - 1)$.

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 . \square

We think of W_1^A and W_0^A as the influence for agent A for distributions that we can construct. Let $I_c^{a,b}$ denote the interval between W_1^A and W_0^A . Note that we do not know which of W_1^A or W_0^A is greater. Claim 3 implies that:

Claim 4. $I_m^{a,b} \cap I_c^{a,b} \neq \emptyset$.

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

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

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 5. 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 , which is the influence to player A when we choose a permutation according to $w^A(a-1, b)$ and append a 'B'. However, since allocating an additional item to player B in any fixed allocation can only degrade player A 's payoff, it must be that $W_0^A \leq w^A(a-1, b)$.

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 $M[a, b]$. 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 implies 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]:

Theorem 6 (Carathodory). *Given a set $V \subset \mathbb{R}^n$ and a point $p \in \text{Conv}V$ — 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. \square

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-player influence maximization problem.

5 A Strategyproof Mechanism for Three or More Players

To construct a strategyproof mechanism for $k > 2$ players, we will impose additional restrictions on the influence functions f_1, \dots, f_k . These additional assumptions are satisfied by many models of influence spread considered in the literature, as we discuss below. 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 mechanism than the mechanism for the case of only two players⁹.

Assumption 1: Mechanism Indifference We will assume that $f(\mathbf{S}) = f(\mathbf{S}')$ whenever the sets $\bigcup_i S_i$ and $\bigcup_i S'_i$ are equal¹⁰. That is, social welfare does not depend on the manner in which allocated items are partitioned between the agents. We will call this the *Mechanism Indifference* (MeI) assumption.

If assumption 1 holds, then we can imagine a greedy algorithm that chooses which items to add to the set $\bigcup_i S_i$ one at a time to greedily maximize the social welfare. By assumption 1, the welfare does not depend on how these items are divided among the players. This greedy algorithm generates a certain social welfare whenever the sum of budgets is t ; write $w(t)$ for this welfare. Note that $w(0), w(1), \dots$ is a concave non-decreasing sequence.

Assumption 2: Agent Indifference We will assume that $f_i(S_i, \mathbf{S}_{-i}) = f_i(S_i, \mathbf{S}'_{-i})$ 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 essentially vacuous.

We note that the models for competitive influence spread proposed by Carnes et al. [7] and Bharathi et al. [3] are based on a cascade model of influence spread, and satisfy both the MeI and AgI assumptions.

⁸Note 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.

⁹At 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 3 show that this is not just a nuance of the proof but rather an intrinsic obstacle to using the uniform distribution.

¹⁰We again postpone discussion of multiset allocations until Appendix A

Similarly, if we restrict the OR model in [5] so that the underlying spread process is a cascade (and not a threshold) process and agents are anonymous (a restriction that will be defined in Section 6), as assumed in the Carnes et al models, then this special case of the OR model also satisfies MeI and AgI.

5.1 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. The MeI assumption implies that the random greedy mechanism obtains a constant factor approximation

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 $u_i \leftarrow \operatorname{argmax}_{c \in U} \{f(I \cup \{c\}) - f(I)\};$
- 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**

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 7. *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 [23, 13], 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 . 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 . 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 7 makes crucial use of the fact that there are at least three players. Indeed, in Appendix B we give an example satisfying the MeI and AgI assumptions for which the random greedy algorithm is not strategyproof for two players.

6 Disjoint Allocations

One feature of the mechanism from Section 4 for 2 players is that it does not necessarily return a profile of disjoint allocations. That is, the algorithm may allocate a given element to both agents. Disjointness is a natural property to require in many models of influence in social networks. Note that the mechanism from Section 5 always returns disjoint allocations.

We now show how to modify the mechanism from Section 4 to ensure disjoint allocations. Recall that our general strategy in the non-disjoint case was to use the locally greedy algorithm and construct a strategyproof-inducing distribution over player orderings for that algorithm. Our strategy for achieving disjointness will be to modify the underlying greedy algorithm so that it only returns disjoint allocations, then apply the same techniques as in Section 4 to convert this algorithm into a strategyproof mechanism. As noted in Section 4, our method can be applied to *any* myopic allocation with a social welfare guarantee that does not depend on the chosen order of players. It therefore suffices to find such a myopic allocation method that guarantees disjointness.

When the disjointness constraint is combined with demand restrictions, the set of valid allocations is not a matroid but rather an intersection of two matroids. The locally greedy algorithm described in Section 2 is therefore not guaranteed to obtain a constant approximation for every ordering of the players. For example, suppose the ground set U consists of two items, 1 and 2. Suppose player A has values 1 and $1 + \epsilon$ for items 1 and 2, respectively (where $\epsilon > 0$ is arbitrarily small), and player B has values 1 and N for items 1 and 2, respectively (where $N > 1$ is arbitrarily large). When the demands of the two players are 1, the locally greedy algorithm might allocate to either player first, but if it allocates to player A first then it obtains the unbounded approximation ratio $\frac{N+1}{2+\epsilon}$.

The above problem stems from the asymmetry in the valuations of the two players. To address this issue, we introduce a notion of player anonymity that captures those circumstances in which these problems do not occur.

Definition 8. We say agents are anonymous if their valuations are symmetric; that is, $f_i(S_1, \dots, S_k) = f_{\pi(i)}(S_{\pi(1)}, \dots, S_{\pi(k)})$ for all permutations π and all agents $1 \leq i \leq k$.

If players are anonymous then the social welfare satisfies $f(S_1, \dots, S_k) = f(S_{\pi(1)}, \dots, S_{\pi(k)})$ for all permutations π . We note that the influence models proposed by Carnes et al. [7] and Bharathi et al. [3] satisfy this condition. At the end of this section we will discuss the relationship between the anonymity condition and the Agent Indifference and Mechanism Indifference conditions from Section 5.

What we now show is that when the players are anonymous, our order-independent locally greedy algorithm from Section 2 obtains a strategyproof mechanism with a $(k+1)$ -approximation to the optimal social welfare, if the given permutation over orderings of the player allocations is sampled from a truthfulness-inducing distribution over permutations (e.g. the distributions we have obtained in the case of two players). Hence, this method provides a transformation to the disjoint allocations case, if one were to obtain a distribution over permutations for the non-disjoint case.

Algorithm 3 is a simple modification to Algorithm 1, in which we explicitly enforce disjointness of the allocations.

Algorithm 3: Disjoint Locally Greedy algorithm

Input: Ground set $U = \{e_1, \dots, e_n\}$, demands a, b for players $1, \dots, k$, a valid permutation

$\pi : \{1, 2, \dots, t\} \rightarrow \{1, 2, \dots, k\}$ where $t = \sum_{i=1}^k b_i$

Output: An allocation $I_1, \dots, I_k \subseteq U$ for the k players

```

1 for  $j = 1 \dots k$  do
2   |  $I_j = \emptyset$ 
3 end
4 for  $i = 1 \dots b_1 + \dots + b_k$  do
5   |  $u \leftarrow \operatorname{argmax}_{c \in U - \cup I_j} \{w(I_{\pi(i)} \cup \{c\}, I_{-i}) - w(I_{\pi(i)}, I_{-i})\}$ ;
6   |  $I_{\pi(i)} \leftarrow I_{\pi(i)} \cup \{u\}$ ;
7 end

```

Theorem 9. For any permutation (of player allocation turns) $\pi : \{1, 2, \dots, t\} \rightarrow \{1, 2, \dots, k\}$ where $t = \sum_{i=1}^k b_i$, Algorithm 3 obtains $(k+1)$ -approximation to the optimal social welfare obtainable for disjoint allocation for identical players $1, \dots, k$.

Proof. Let $\mathbf{O} = (O_1, \dots, O_k)$ be an optimal allocation. Let (I_1, \dots, I_k) be the allocation obtained by running Algorithm 3 for some permutation π . Partition \mathbf{O} as follows. For each player i , set $O_i^j = O_i \cap I_j$ for all $j \neq i$, and let $O_i^0 = O_i - \bigcup_{j \neq i} I_j$. By submodularity,

$$w(O_1, \dots, O_k) \leq w(O_1^0, \dots, O_k^0) + \sum_{i=1}^{k-1} w(O_1^{(1+i) \bmod k}, \dots, O_k^{(k+i) \bmod k}). \quad (4)$$

Due to anonymity and the fact that $O_i^j \subseteq I_j$, for all $j \in [k]$ we get

$$\sum_{i=1}^{k-1} w(O_1^{(1+i) \bmod k}, \dots, O_k^{(k+i) \bmod k}) \leq (k-1) \cdot w(I_1, \dots, I_k). \quad (5)$$

Next, we can apply the analysis performed for the original locally greedy algorithm, so as to obtain the relation specified in the following lemma

Lemma 10. $w(O_1^0, \dots, O_k^0) \leq 2 \cdot w(I_1, \dots, I_k)$.

For readability, we prove the lemma at the end of this section. Now, combining relations (4) and (5) and applying Lemma 10, we get:

$$w(O_1, \dots, O_k) \leq 2 \cdot w(I_1, \dots, I_k) + (k-1)w(I_1, \dots, I_k) = (k+1) \cdot w(I_1, \dots, I_k) \quad (6)$$

□

Corollary 11. *For $k = 2$ players, there exists a 3-approximate strategyproof mechanism that achieves disjoint allocations.*

Proof. Observe that this revised version of the locally greedy algorithm is order-independent. That is, we obtain the same (constant) bound on its approximation ratio for any player ordering. In particular, this means that we can apply the mechanism described in Section 4 for obtaining a strategyproof solution without significantly degrading the approximation ratio of the greedy algorithm. \square

We note that for all k there is a natural greedy algorithm for this problem (with disjointness) that obtains a 3-approximation for any k . Namely, the greedy algorithm that chooses both the player and the allocation that maximizes the marginal utility on each iteration [22]). However, this algorithm imposes a particular ordering on the allocations and therefore does not allow the degree of freedom required by our truthful mechanism construction.

Proof of Lemma 10. We now adapt the analysis of the locally greedy algorithm (e.g. [13]) in order to prove the bound in Lemma 10. We begin by introducing some additional notation. First, for $i \in [k]$, let $O'_i = O_i^0 \setminus I_i$. For an item $e \in I_i$ ($i \in [k]$), $\mathbf{S}^e = (S_1^e, \dots, S_k^e)$ denotes the partial solution of the algorithm at the time of item e 's addition. For an allocation $(A_1, \dots, A_k) \subseteq U^k$ and an item $e \in U$, define the marginal gain to be $\rho_e^i(A_1, \dots, A_k) = w(A_1, \dots, A_i \cup \{e\}, \dots, A_k) - w(A_1, \dots, A_k)$. Lastly, for $i \in [k]$, we let $e_i = \arg \min_{e \in I_i} \rho_e^i(\mathbf{S}^e)$; i.e. the minimal marginal increase to social welfare, obtained by the algorithm, when adding an element to I_i . The following lemma follows from the fact that the social welfare is a non-decreasing, submodular function:

Lemma 12 ([23]).

$$w(O_1^0, \dots, O_k^0) \leq w(I_1, \dots, I_k) + \sum_{i=1}^k \sum_{e \in O'_i} \rho_e^i(I_1, \dots, I_k)$$

Now, by the greedy rule of Algorithm 3, we have:

$$\rho_{e_i}^i(S_i^e) \geq \rho_e^i(S_i^e), \text{ for all } i \in [k], \text{ and for all } e \in U \setminus S_i^{e_i} \quad (7)$$

Additionally, using the submodularity of $w(\cdot)$, and the fact that for all $i, j \in [k]$, $S_j^{e_i} \subseteq I_j$, we get:

$$\begin{aligned} w(O_1^0, \dots, O_k^0) &\leq w(I_1, \dots, I_k) + \sum_{i=1}^k \sum_{e \in O'_i} \rho_e^i(I_1, \dots, I_k) \\ &\leq w(I_1, \dots, I_k) + \sum_{i=1}^k \sum_{e \in O'_i} \rho_{e_i}^i(\mathbf{S}^{e_i}) \end{aligned} \quad (8)$$

Now, using (7), we further extend the above bound as follows

$$\begin{aligned} w(O_1^0, \dots, O_k^0) &\leq w(I_1, \dots, I_k) + \sum_{i=1}^k \sum_{e \in O'_i} \rho_{e_i}^i(\mathbf{S}^{e_i}) \\ &= w(I_1, \dots, I_k) + \sum_{i=1}^k b_i \cdot \rho_{e_i}^i(\mathbf{S}^{e_i}) \\ &\leq w(I_1, \dots, I_k) + w(I_1, \dots, I_k) = 2 \cdot w(I_1, \dots, I_k) \end{aligned} \quad (9)$$

where the last inequality follows from the definition of the elements e_1, \dots, e_k . \square

6.1 Relation to Indifference Conditions

In this section we explore the relationship between the anonymity condition required by Theorem 9 and the mechanism and agent indifference conditions (MeI and AgI) used in Section 5. As we will show, these conditions are incomparable when there are only two players, but when there are three or more players the AgI and MeI conditions together are strictly stronger than the anonymity condition.

Consider first the case of two players. To see that MeI does not imply anonymity, consider the following example with two objects $\{a, b\}$ and two players. The functions f_1 and f_2 are given by $f_1(x, 0) = f_2(0, x) = 2$ for any singleton x , $f_1(\{a, b\}, 0) = f_2(0, \{a, b\}) = 3$, and $f_1(x, y) = 1.6$, $f_2(x, y) = 1.4$ for $(x, y) = (a, b)$ or vice-versa. One can verify that $f = f_1 + f_2$ is submodular and that adverse competition and mechanism indifference are satisfied, but it is not anonymous (since $f_1(x, y) \neq f_2(y, x)$ for singletons x and y).

To see that anonymity does not imply MeI, consider the following example with two objects $\{a, b\}$ and two players. We will have $f_1(x, 0) = f_2(0, x) = 1$ for each singleton x , $f_1(\{a, b\}, 0) = f_2(0, \{a, b\}) = 2$, but $f_1(x, y) = f_2(x, y) = 3/4$ for $(x, y) = (a, b)$ or vice-versa. This pair of functions exhibits adverse competition and its sum is submodular, but it does not satisfy MeI (since $f(a, b) \neq f(\{a, b\}, 0)$).

For $k \geq 3$ players, MeI and AgI together imply anonymity.

Theorem 13. *If there are $k \geq 3$ agents and the AgI and MeI conditions hold, then the agents are anonymous.*

Proof. We will assume $k = 3$ for notational convenience; extending to $k > 3$ is straightforward. It is sufficient to show that $f_1(S, T, U) = f_2(T, S, U)$ for arbitrary sets S, T, U ; symmetry with respect to all other permutations then follows by composing transpositions.

We first show that $f_1(S, \emptyset, U) = f_2(\emptyset, S, U)$ for all sets S and U . By MeI, $f(S, \emptyset, U) = f(\emptyset, S, U)$. By AgI, $f_3(S, \emptyset, U) = f_3(\emptyset, S, U)$. By our assumed normalization, $f_2(S, \emptyset, U) = f_1(\emptyset, S, U) = 0$. Taking sums, we conclude that $f_1(S, \emptyset, U) = f_2(\emptyset, S, U)$.

By the same argument, $f_1(S, \emptyset, U \cup T) = f_2(\emptyset, S, U \cup T)$. But then, by AgI, $f_1(S, T, U) = f_1(S, \emptyset, U + T) = f_2(\emptyset, S, U + T) = f_2(T, S, U)$, as required. \square

Finally, we show that the MeI and AgI assumptions together are strictly stronger than anonymity for $k \geq 3$ players, as anonymity does not imply MeI. Consider the following example with 3 objects $\{a_1, a_2, a_3\}$ and 3 players. For any labeling of the singletons as x, y, z , define $f_1(x, y, z) = 7/24$, $f_1(\{x, y\}, z, 0) = f_1(\{x, y\}, 0, z) = 3/4$, $f_1(x, \{y, z\}, 0) = f_1(x, 0, \{y, z\}) = 1/4$, and $f_1(\{x, y, z\}, 0, 0) = 1$. Define f_2 and f_3 symmetrically, so agents are anonymous. Adverse competition is satisfied and the sum of these functions is submodular, but neither MeI nor AgI are satisfied.

7 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.

An important question is how to extend our results to the more general case of $k > 2$ agents without the MeI and AgI assumptions. As discussed in Appendix D, 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 [27] or multiple [19] 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. Without this 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.

An important question is how to extend our results to the more general case of k players. As we discuss in Appendix D, it seems that a fundamentally new approach would be required to obtain a polytime strategyproof mechanism for $k > 2$ players.

Acknowledgements We would like to thank Yuval Filmus for many helpful comments and discussions.

References

- [1] K. Apt and E. Markakis. Diffusion in social networks with competing products. In G. Persiano, editor, *Algorithmic Game Theory*, volume 6982 of *Lecture Notes in Computer Science*, pages 212–223. Springer Berlin / Heidelberg, 2011.
- [2] I. Ashlagi, F. Fischer, I. Kash, and A. D. Procaccia. Mix and match. In *Proceedings of the 11th ACM conference on Electronic commerce*, EC '10, pages 305–314, New York, NY, USA, 2010. ACM.
- [3] S. Bharathi, D. Kempe, and M. Salek. Competitive influence maximization in social networks. In X. Deng and F. Graham, editors, *Internet and Network Economics*, volume 4858 of *Lecture Notes in Computer Science*, pages 306–311. Springer Berlin / Heidelberg, 2007.
- [4] A. Borodin, M. Braverman, B. Lucier, and J. Oren. Truthful mechanisms for competing submodular processes. *CoRR*, abs/1202.2097, 2012.
- [5] A. Borodin, Y. Filmus, and J. Oren. Threshold models for competitive influence in social networks. In A. Saberi, editor, *Internet and Network Economics*, volume 6484 of *Lecture Notes in Computer Science*, pages 539–550. Springer Berlin / Heidelberg, 2010.
- [6] J. J. Brown and P. H. Reingen. Social ties and word of mouth referral behavior. *Journal of Consumer Research*, 14(3):350–362, 1987.
- [7] T. Carnes, C. Nagarajan, S. M. Wild, and A. van Zuylen. Maximizing influence in a competitive social network: a follower’s perspective. In *Proceedings of the ninth international conference on Electronic commerce*, ICEC '07, pages 351–360, New York, NY, USA, 2007. ACM.
- [8] D. Centola and M. Macy. Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3):702–734, 2007.
- [9] P. Dubey, R. Garg, and B. De Meyer. Competing for customers in a social network: The quasi-linear case. In P. Spirakis, M. Mavronicolas, and S. Kontogiannis, editors, *Internet and Network Economics*, volume 4286 of *Lecture Notes in Computer Science*, pages 162–173. Springer Berlin / Heidelberg, 2006.
- [10] J. Feldman, M. Henzinger, N. Korula, V. S. Mirrokni, and C. Stein. Online stochastic packing applied to display ad allocation. In *Proceedings of the 18th annual European conference on Algorithms: Part I*, ESA'10, pages 182–194, Berlin, Heidelberg, 2010. Springer-Verlag.
- [11] J. Feldman, A. Mehta, V. Mirrokni, and S. Muthukrishnan. Online stochastic matching: Beating $1-1/e$. *FOCS*, 0:117–126, 2009.
- [12] J. Goldenberg, B. Libai, and E. Mulle. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing Letters*, pages 221–223, 2001.
- [13] P. Goundan and A. Schulz. Revisiting the greedy approach to submodular set function maximization. *Optimization online*, (1984):1–25, 2007.
- [14] S. Goyal and M. Kearns. Competitive contagion in networks. In *STOC*, pages 759–774, 2012.

- [15] M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, (83):1420–1443, 1978.
- [16] D. Kempe, J. Kleinberg, and E. Tardos. 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*, pages 137–146, New York, NY, USA, 2003. ACM.
- [17] D. Kempe, J. Kleinberg, and É. Tardos. Influential nodes in a diffusion model for social networks. In L. Caires, G. Italiano, L. Monteiro, C. Palamidessi, and M. Yung, editors, *Automata, Languages and Programming*, volume 3580 of *Lecture Notes in Computer Science*, pages 1127–1138. Springer Berlin / Heidelberg, 2005.
- [18] J. Kostka, Y. A. Oswald, and R. Wattenhofer. Word of mouth: Rumor dissemination in social networks. In *Proceedings of the 15th international colloquium on Structural Information and Communication Complexity*, SIROCCO '08, pages 185–196, Berlin, Heidelberg, 2008. Springer-Verlag.
- [19] A. Kulik, H. Shachnai, and T. Tamir. Maximizing submodular set functions subject to multiple linear constraints. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '09*, pages 545–554, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.
- [20] P. Lu, X. Sun, Y. Wang, and Z. A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM conference on Electronic commerce, EC '10*, pages 315–324, New York, NY, USA, 2010. ACM.
- [21] E. Mossel and S. Roch. Submodularity of influence in social networks: From local to global. *SIAM J. Comput.*, 39(6):2176–2188, 2010.
- [22] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions—II. *Mathematical Programming*, 14(1):265–294, Dec. 1978.
- [23] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14:265–294, 1978. 10.1007/BF01588971.
- [24] R. T. Rockafellar. *Convex Analysis (Princeton Landmarks in Mathematics and Physics)*. Princeton University Press, December 1996.
- [25] T. Schelling. *Micromotives and macrobehavior*. Norton, 1978.
- [26] Y. Singer. How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In *WSDM*, pages 733–742, 2012.
- [27] M. Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1):41 – 43, 2004.

A 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. ([16]) 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 [17]. In the case of the competitive influence spread models in [3], [7], and [5], 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 ([3], [7], and [5]) 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 absense 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 [7], the propagation of influence upholds both the mechanism and agent indifference properties. In [14], it is assumed that the probability that a node will adopt some technology is a function of the fraction of influenced neighbours (regardless of their assumed technology). This immediately implies mechanism indifference, as general spread is invariant with respect to the distribution of technologies among initial nodes. For their positive price of anarchy results about more than two players, it is assumed that the selection function is linear which would imply mechanism indifference.

Anonymity With the exception of the OR model ([5]), the above mentioned models also satisfy an anonymity assumption that will be needed to modify the local greedy algorithm (as in Algorithms 1 and 2) so as to insure that the initial allocation is disjoint (see Appendix 6). Anonymity basically means that the players are symmetric and when there are more than two players this is a somewhat weaker condition than having both mechanism and agent indifference as we shall show in Appendix 6. We note that in [3] and [7] there is only one edge-weight per edge ¹¹ thereby enforcing anonymity. In [14], it is explicitly stated that the selection function is symmetric across the players and this implies anonymity.

¹¹In fact, towards the end of the paper, the authors of [7] conjecture that their results extend to the non-anonymous case where each edge has technology-specific weights. This conjecture was later shown to be false in [5].

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 [14]. 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 2-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.


B Counter examples 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 [22] (see also Goundan and Schultz [13]) 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).

B.1 The OR model

We will consider examples of a special case of the OR model for influence spread, as studied in [5]. 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 C we show how to modify the examples given in this section to be unweighted. The process then 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 $w(u_A, v_A)$. 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. [16]. Finally, nodes that are contained in $O_A \setminus O_B$ will be assigned to player A , nodes in $O_B \setminus O_A$



(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.

Figure 1: Counter-examples for the mechanism under the deterministic dictatorship and Round Robin orderings. In both case, we set the weights $w_{c_i} = \epsilon$ and $w_{u_i} = 1$, for all $1 \leq i \leq 4$. $0 < \epsilon < \frac{1}{8}$

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.

B.2 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 Figure 1a. 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_4 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 $\frac{1}{10} + \frac{1}{2} \cdot \frac{9}{10} = \frac{11}{20}$. 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 Figure 1b. 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 .

B.3 The uniform random greedy algorithm is not strategyproof

As we shall see in Section 5, 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 $\frac{\epsilon}{\epsilon-1}$ 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




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

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 Figure 2. We note that for this example, Algorithm 2 in Section 5 is equivalent to first choosing a random order of allocation (e.g. choosing all possible permutations satisfying agent demands 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 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.

C Counterexamples with Unweighted Nodes

In Section 3 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 Section B.3 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 $2/\epsilon$ 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 Section B.3, to demonstrate that if agent B declares 1 then agent A would rather declare 3 than 4.

D Tightness of Approach: More than Two Players

The mechanism we construct in Section 4 is applicable to settings in which there are precisely two competing players, and our mechanism in Section 5 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 Section 4, 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 k , then we can extend this to a strategyproof mechanism for all possible budget declarations that total at most $k + 1$. One key observation that makes this extension possible is that the strategyproofness condition can be re-expressed as a certain “adverse competition” property: if one player increases his budget, then the expected utility for the other player cannot increase by more than the marginal gain in total welfare. In other words, in the notation of Lemma 1, we can construct our mechanism so that for all $a + b \geq 1$, $w^A(a, b) \leq w^A(a, b - 1) + \Delta^{\oplus B}(a, b)$.

A direct extension of our approach to three players would involve proving that an allocation rule that is strategyproof and satisfies the adverse competition condition for all budgets that total at most k can always be extended to handle budgets that total up to $k + 1$. 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 c of value; all other elements are worth nothing. The utility for each agent is 1 if their allocation contains c , otherwise their utility is 0. In this case, the locally greedy algorithm simply gives element c 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 adverse competition property, and also satisfies the cross-monotonicity property (i.e. the first invariant of Lemma 1). However, we claim that no allocation on input $(1, 1, 1)$ that obtains positive social welfare can maintain the adverse competition property. To see this, note that the adverse 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 adverse 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 mechanism for budgets totalling at most 2 to a (strategyproof, cross-monotone) 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 a decrease in the utility of each other 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. Hence, we believe that our greedy mechanism cannot be made strategyproof via our general 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. [20] results for k -facility location. Additionally Ashlagi et al. discussed similar issues ([2]) 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 $O(n)$.