

Multi-Winner Contests for Strategic Diffusion in Social Networks

Wen Shen, Yang Feng, Cristina V. Lopes

University of California, Irvine, California 92697, United States
wen.shen@uci.edu, yang.feng@uci.edu, lopes@ics.uci.edu

Abstract

Strategic diffusion encourages participants to take active roles in promoting stakeholders' agendas by rewarding successful referrals. As social media continues to transform the way people communicate, strategic diffusion has become a powerful tool for stakeholders to influence people's decisions or behaviors for desired objectives. Existing reward mechanisms for strategic diffusion are usually either vulnerable to false-name attacks or not individually rational for participants that have made successful referrals. Here, we introduce a novel *multi-winner contests* (MWC) mechanism for strategic diffusion in social networks. The MWC mechanism satisfies several desirable properties, including false-name-proofness, individual rationality, budget constraint, monotonicity, and sub-graph constraint. Numerical experiments on four real-world social network datasets demonstrate that stakeholders can significantly boost participants' aggregated efforts with proper design of competitions. Our work sheds light on how to design manipulation-resistant mechanisms with appropriate contests.

Introduction

Strategic diffusion is the process of spreading information among social media users to promote desired private or social outcomes (Galeotti and Goyal 2009). As the impact of social media on people's daily lives continues to grow, strategic diffusion has become a prominent tool for stakeholders (e.g., individuals, companies, governments, and NGOs) to influence people's preferences, decisions or behaviors (Jackson and Yariv 2011; Chaffey 2016). The vast popularity of strategic diffusion in social networks is primarily because it encourages participants to take active roles in promoting stakeholders' agendas in a word-of-mouth fashion (i.e., in the forms of referrals). This viral marketing strategy can reach a broader audience at a faster pace and are usually more economically efficient than traditional advertising such as newspapers, radios or televisions (Leskovec, Adamic, and Huberman 2007; Galeotti and Goyal 2009).

To unleash the power of crowds, stakeholders usually reward both direct and indirect referrals to encourage potential participants to perform the tasks early and invite influential players to participate. For example, if Alice refers Bob and

Bob then refers Cathy, both Alice and Bob are rewarded for Cathy's purchases. However, Bob is typically given more rewards than Alice for his direct referral. Such referral mechanisms are often called the *geometric reward mechanisms*, or the *incentive tree mechanisms* (Pickard et al. 2011). Geometric reward mechanisms are usually effective and easy to implement (Tang et al. 2011). As a result, they have witnessed a growing range of serious applications, including product promotion (Drucker and Fleischer 2012), disaster rescue (Rahwan et al. 2013), global manhunt (Rahwan et al. 2013), participatory sensing (Gao et al. 2015) and crowd-funding (Naroditskiy et al. 2014b).

Despite the promising prospects of the geometric reward mechanisms, they are usually vulnerable to manipulations (Drucker and Fleischer 2012). Under many geometric reward mechanisms, indirect referrals are rewarded (Pickard et al. 2011). A strategic player may create multiple fake accounts or identities on his behalf with one referring another to increase his rewards. Players' such malicious behaviors are often called *false-name attacks*. A false-name attack refers to a strategy that a profit-maximizing agent utilizes to gain benefits by creating multiple false identities (i.e., replicas). False identities can be analogous to "free riders" in the sense that they pay less than what they should have paid. This free rider problem arises when genuine identities are unfeasible or difficult to be recognized. Unfortunately, many social network services lack effective methods to fully eliminate false identities (Ferrara et al. 2016).

False-name attacks severely impede stakeholders from implementing desired individual or societal outcomes. On the one hand, false-name attacks are undesirable because they not only diminish stakeholders' revenues but also reduce other truthful players' payoffs (Drucker and Fleischer 2012). On the other hand, prior research indicates that false-name attacks are pervasive in social networks since players may create false accounts with no or minimal efforts if no interventions are given (Lorenz et al. 2011; Naroditskiy et al. 2014a; Ferrara et al. 2016).

Despite that much research has been devoted to tackling the false-name-attack problems in social networks (Conitzer et al. 2010; Emek et al. 2011; Todo, Iwasaki, and Yokoo 2011; Brill et al. 2016), our work is most closely related to the mechanism design problem for multi-level marketing (Drucker and Fleischer 2012). In their work, Drucker

and Fleischer introduce a class of mechanisms by capping the rewards a player can get from indirect referrals. They show that under their mechanisms false-name attacks are unprofitable. While illuminating, their methods indicate that some of the less influential players in the referral networks will receive no rewards for successful indirect referrals. For instance, players with ordinary or low capabilities may have little or no incentives to participate in the referral mechanisms. Ignoring these players is problematic for stakeholders because the population of many online platforms typically consist of a large portion of participants with low or common capabilities (Ipeirotis 2010; Dow et al. 2012). Besides, the implementation of their mechanisms is not based on graphs but rather on trees, it remains unknown whether their mechanisms can be applied to large-scale social networks consisting of thousands of participants with different abilities.

To address these issues, we introduce a novel mechanism called the *Multi-Winner Contests* (MWC) mechanism for strategic diffusion in social networks. The MWC mechanism distinguishes itself from existing methods in two aspects. First, it allocates virtual credits for each successful referrals by taking both diffusion contributions and verifiable task efforts into account. Second, it determines the diffusion rewards for each player according to the results of contests that compare the virtual credits earned by the qualified players. We show that the mechanism is false-name-proof, individually rational, budget-constrained, monotonic, subgraph-constrained, and computationally efficient. We conducted extensive experiments with four real-world social network datasets. Experimental results demonstrate that false-name-attacks are unprofitable under the MWC mechanism. Stakeholders can significantly boost the aggregated efforts of players when they select parameters of the MWC mechanism appropriately. Our work casts light on how to integrate competitions into the design of novel mechanisms to counter manipulations.

Strategic Diffusion in Social Networks

This section first describes notations used for modeling strategic diffusion in social networks. It then introduces the mechanism design problem for strategic diffusion. After defining the concept of false-name attacks, it presents the solution concepts for the reward mechanism design problem.

Preliminaries

Strategic diffusion processes or referral networks are usually modeled with *directed acyclic graphs* (DAGs). We consider a referral DAG $G = (V, E)$ (See Figure 1) where V denotes the set of players that may contribute (i.e., nodes) and E denotes the set of referral relationships (i.e., edges). For any nodes $v, u \in G$ ($v \neq u$), if there is a directed edge from v to u , then it means that u 's decision to contribute is partially a result of v . In this case, we say u is a *direct successor* of v and v is a *direct predecessor* of u . For each node v , the number of direct predecessors v has is its *indegree* $deg^-(v)$. The number of direct successors it has is its *outdegree*, denoted by $deg^+(v)$. A source node (i.e., a seed node) has an indegree of 0, while a sink node has an outdegree of 0.




Figure 1: A typical directed acyclic graph.

For any nodes $v, u \in G$ ($v \neq u$), if a path leads from v to u , then u is said to be a *successor* of v and reachable from u , and v is said to be a *predecessor* of v . We write the distance between v and u by $dist(u, v)$. If there is no path from v to u , then the distance between the two nodes is infinity; that is $dist(v, u) = \infty$. The length from v to v is 0, i.e., $dist(v, v) = 0$. Let κ_v^+ be the set of all successors of v , we have $\kappa_v^+ = \{u \in V : 0 < dist(v, u) < \infty\}$. Similarly, the set of v 's predecessors $\kappa_v^- = \{u \in V : 0 < dist(u, v) < \infty\}$. Let G_v be the subgraph rooted at v , then $G_v = (V', E')$, where $V' = \{v\} \cup \kappa_v^+$ and $E' \subseteq E$.

In a DAG $G = (V, E)$, if a player $v \in V$ makes a contribution to the designated tasks (e.g., making purchases, answering questions, reporting software bugs), we say player v exerts *task efforts* $t_v \in \mathbb{R}_{\geq 0}$. Player v may either directly or indirectly spread the information of the tasks to his successors to maximize his profits. If one of his successors u also contributes to the tasks, then we say v makes *diffusion contributions* $d_v \in \mathbb{R}_{\geq 0}$. In our model, we assume that the task efforts are verifiable and available for the stakeholder. For instance, a seller usually has the legitimate transactions information about the products she sells to each buyer.

Each edge $e = (v_i, v_j) \in E$ has a weight $\omega_{v_i v_j}$ (i.e., ω_e). The weight $\omega_{v_i v_j}$ represents the proportion of the credits assigned to v_i when v_j contributes to the designated tasks with efforts t_{v_j} . Here, $\omega_{v_i v_j} = 1/deg_{v_j}^-$ if $v_i \neq v_j$; otherwise, $\omega_{v_i v_j} = 1$. Let $P_{v_i v_j}$ be the set of all paths from v_i to v_j and p be a path in the set $P_{v_i v_j}$. Path p 's weight $\omega(p)$ is the product of the weights of all edges along the path. That is, $\omega(p) = \prod_{(v_i, v_j) \in p} \omega_{v_i v_j}$. Here, the length of the path $|p| = dist(v_i, v_j)$.

Mechanism Design for Strategic Diffusion

We consider a principal (e.g., a seller, a task owner) employs strategic diffusion to maximize the aggregated efforts on the tasks she designates. Initially, the principal selects some participants (i.e., seeds) $S \subseteq V$ in a social network $G = (V, E)$ to perform the tasks that she specifies. The principal may select the seed nodes randomly if she has no prior knowledge of the social network. Alternatively, she may select the seed nodes using the influence maximization approach if the network structure and diffusion method are known a priori (Kempe, Kleinberg, and Tardos 2003).

Some of the seed participants then perform the tasks and invite their neighbors in the social network to participate. If a participant v exerts task efforts t_v , he will receive a task reward $\pi_t(v)$ from the principal. Participant v may spread the information to his successors $u \in \kappa_v^+$ to maximize his profits. If a successor u exerts task efforts t_u , the referrer v will receive a diffusion reward $\pi_d(v, u)$. The total diffusion rewards for v : $\pi_d(v) = \sum_{u \in \kappa_v^+} \pi_d(v, u)$. Thus, the total re-

wards for v are determined by: $\pi(v) = \pi_t(v) + \pi_d(v)$.

The total contributions of v include both the task efforts and the diffusion contributions. In strategic diffusion, the task efforts are verifiable while the diffusion contributions are difficult to verify because strategic players can generate fake referrals with no or minimal efforts by creating multiple false identities with one being referred by another. We thus treat them separately. Let $c(v)$ be the total contributions for v , we have: $c(v) = (t_v, d_v)$. Let Θ_v be the space of v 's total contributions $c(v)$, and $\Theta = (\Theta_v)_{v \in V}$.

In our model, we assume that the valuation of a player is linear in the rewards he receives and the cost he pays is linear in the task efforts he has contributed. That is, there are constants $\alpha_v, \beta_v, \gamma_v$ and ς_v with $\beta_v > 0$ and $\varsigma_v > 0$, such that player v 's utility is $U(v) := \alpha_v + \beta_v \cdot \pi(v) - (\gamma_v + \varsigma_v \cdot t_v)$. Without loss of generality, let $\alpha_v = \gamma_v$ and $\delta_v = \frac{\varsigma_v}{\beta_v}$. We have player v 's utility:

$$U(v) = \pi(v) - \delta_v \cdot t_v, \quad (1)$$

where $\delta_v > 0$ is a private coefficient that determines the player's marginal cost for exerting extra unit effort. The higher ability a player owns, the lower δ_v he has, and vice versa. That is, there is a negative correlation between a player's ability and his cost coefficient.

We note that a player may incur costs on spreading the information to his neighbors. We avoid explicitly including players' diffusion costs into the utility function for two concerns. First, diffusion efforts are not directly verifiable for the principal. The exact correlations between the costs on task efforts and the diffusion costs may vary from player to player, which are unknown to the principal. Second, a player v can integrate his diffusion costs into the task efforts by setting a higher cost coefficient δ_v . Similar techniques were also used in literature (Shen et al. 2018).

In strategic diffusion, the principal is interested in a *reward mechanism* π that determines the reward for each player that has exerted efforts.

Definition 1. A reward mechanism π is a tuple of payments for each player $v \in V$, where $G = (V, E)$. That is, $\pi = (\pi(v))_{v \in V}$, where $\pi(v) : \Theta \rightarrow \mathbb{R}$.

False-Name Attacks

In a graph $G = (V, E)$, we say a graph $G' = (V'', E'')$ and a set of replicas $R \subseteq V'$ are a *false-name attack* by v in G if when we collapse R into the single node with label v in G' we get the graph G .

Definition 2 (False-Name Attack). Given a referral graph $G = (V, E)$, for any $v \in V$, let DS_v, DP_v be the sets of node v 's direct successors and direct predecessors, respectively. $G' = (V'', E'')$ is obtained from G by a false-name attack at v if:

- $V'' = V \setminus \{v\} \cup \{r_1, \dots, r_m\}$, $m > 1$. The set of nodes $R = \{r_1, \dots, r_m\}$ is the set of replicas of v .
- The sum of the task efforts of v 's replicas is equal to v 's task efforts. That is, $t_v = \sum_{r \in R} t_r$, where $t_r > 0$.
- All replicas of v have at least one direct successor; that is $\deg^+(r) \geq 1$ for all $r \in R$.




Figure 2: False-name attacks at node 2: nodes $2_1, 2_2, \dots, 2_m$ are replicas of node 2.

- For all $u \in DS_v$, $G_u = G'_u$, where G'_u is a subgraph of G' rooted at u .
- The direct predecessors of a replica of v are either replicas of v , or the direct predecessors of $v \in V$, or null. That is, $DP_r = R \setminus \{r\} \cup DP_v \cup \emptyset$, where $r \in R$.
- $\forall r \in R, \exists$ a replica r' of v such that $r \in \kappa_{r'}^+$.

In general, there are three types of false-name attacks: type 1 (See Figure 2a), type 2 (See Figure 2b), and hybrid. Type 1 false-name attacks occur in the form of a long referral chain that consists of the replicas, while in type 2 attacks the replicas operate in parallel. Hybrid attacks are combinations of type 1 and type 2 attacks.

A false-name attack is profitable if the sum of the rewards received by the replicas $r \in R$ of $v \in V$ are higher than the reward v receives. That is, $\sum_{r \in R} \pi(r) > \pi(v)$.

Solution Concepts

Since false-name attacks are harmful to both the stakeholders and other truthful players, it is desirable that the reward mechanism satisfies the *false-name-proofness* property. A reward mechanism is false-name proof if false-name attacks are unprofitable for every player in a social network. That is, the total rewards that $v \in G$ receives are at least the same as the sum of the rewards received by the replicas of v in a graph obtained from a split of G at node v .

Definition 3 (False-Name-Proofness). A reward mechanism π is false-name-proof if for all $v \in G$: $\pi(v) \geq \sum_{r \in R} \pi(r)$, where R is the set of replicas due to a false-name attack at node v .

According to Equations 1 and 4, if player v 's costs on task efforts (i.e., $\delta_v \cdot t_v$) are equal to or higher than the rewards $\pi_t(v)$, then player v will have no incentives to participate in the strategic diffusion without positive diffusion rewards. To encourage potential players to participate, a reward mechanism needs to ensure that a player who has made successful referrals receives positive (expected) rewards. Otherwise, it would be not appealing to players with high marginal costs (i.e., high δ_v). A reward mechanism satisfies *individual rationality* if its (expected) utility is positive and it allocates positive (expected) diffusion rewards to each player that have made successful referrals.

Definition 4 (Individual Rationality). A reward mechanism π is individually rational if for each player $v \in V$ with $d_v > 0$, we have: $U(v) > 0$, and $\pi_d(v, u) > 0$, where $u \in \kappa_v^+$, and $U(v)$ is determined by Equation 1.

In practice, a reward mechanism should be budget constrained. If not, the mechanism will be economically unfeasible for deployment. In our setting, the total rewards distributed to the players should not exceed a fixed portion ϑ of the total aggregated task efforts.

Definition 5 (Budget Constraint). *A reward mechanism π is budget constrained if: $\sum_{v \in G} \pi(v) \leq \vartheta \cdot \sum_{v \in G} t_v$, where ϑ is a positive constant.*

Another important constraint is that a reward mechanism should limit indirect referrals to restrict the formation of long referral chains. This property is usually called the *monotonicity* constraint. A reward mechanism is monotonic if direct referrals receive higher diffusion rewards than or at least the same rewards as indirect referrals. The monotonicity property encourages participants to form short diffusion chains by offering more rewards. As a result, it limits the scope of indirect rewards, which is desirable in practice because many successful referral chains are usually short (Leskovec, Adamic, and Huberman 2007).

Definition 6 (Monotonicity). *A reward mechanism π is monotonic if v_2 is a successor of v_1 , adding a direct successor v_i to v_2 increases v_1 's diffusion rewards $\pi_d(v_1')$ at least as much as the diffusion rewards $\pi_d(v_1'')$ by adding a direct successor v_j to a successor of v_2 , where $t_{v_i} = t_{v_j}$, i.e., $\pi_d(v_1') \geq \pi_d(v_1'')$.*

In strategic diffusion, players should not have incentives to delay performing the tasks to wait for a referral with a more rewarding position in a social network G . To satisfy this constraint, a reward mechanism should determine the rewards for each player $v \in V$ based on the subgraph G_v rooted at v .

Definition 7 (Subgraph Constraint). *A reward mechanism π is subgraph-constrained if $\pi(v)$ only depends on the rooted subgraph G_v .*

Multi-Winner Contests Mechanism

We present a novel reward mechanism called the *Multi-Winner Contests* (MWC) mechanism. It has two key ingredients. The mechanism first calculates the virtual credits for the diffusion contributions of each player with successful referrals. It then determines the diffusion rewards by holding a contest among players that are in his rooted subgraph. The MWC mechanism is computationally efficient and satisfies several desirable properties, including false-name-proofness, individual rationality, budget constraint, monotonicity, and subgraph constraint.

Virtual Credits for Diffusion Contributions

For each newly joined player v that has exerted task efforts t_v , the MWC mechanism pays $\pi_t(v) = \mu \cdot t_v$ ($\mu > 0$) for his task contributions and allocates virtual credits $\eta \cdot (t_v)^2$ for his diffusion contributions. If v has either directly or indirectly referred player $u \in \kappa_v^+$ to participate ($t_u > 0$), his virtual credits b_v are computed by:

$$b_v = \eta \cdot (t_v)^2 + t_v \cdot \sum_{u \in \kappa_v^+} \sum_{p \in P_{vu}} t_u \cdot \omega(p) \cdot \lambda^{|p|}, \quad (2)$$

where $0 < \lambda < 1$, $\eta \geq \lambda/2$, and P_{vu} is the set of paths from v to u .

The initial allowance of virtual credits $\eta \cdot (t_v)^2$ serves dual purposes. First, it ensures that each player has a positive amount of virtual credits to enter into the contests and receives a positive share of diffusion rewards. Second, it allows the mechanism to significantly reduce a player's virtual credits if he splits his task efforts. This dedicated design is an essential step that makes false-name attacks unprofitable. Note that $t_u \cdot \lambda^{|p|}$ is a typical *tree incentive mechanism* (i.e., geometric reward mechanism). The MWC mechanism extends it to a reward mechanism that applies to social networks by introducing the weights $w(p)$. To incentivize players to exert higher task efforts, the MWC mechanism amplifies the virtual credits for diffusion contributions by multiplying the player's task efforts t_v .

Diffusion Rewards

The MWC mechanism performs multi-winner contests to decide the diffusion rewards. In the contests, the winners are chosen simultaneously by a *contest success function* (CSF) that compares the virtual credits of each player. A CSF determines each player's probability of winning the contest in terms of all players' efforts (Skaperdas 1996).

In general, there are two types of contest success functions: the *ratio* form and the *difference* form (Skaperdas 1996). In the ratio form, the winning probabilities depend on the ratio of efforts exerted by each player. In the difference form, they are determined by the difference in efforts that each player has exerted. The ratio form can be naturally applied to large contests that consist of many players, while extending the difference form to contests with more than two players is usually non-trivial and difficult (Jia, Skaperdas, and Vaidya 2013). Strategic diffusion in social networks typically involves a large number of players. It is thus desirable to use the ratio form CSFs for analytical convenience.

The ratio form CSFs typically predict contest outcomes from the ratio of the efforts that each player has devoted. For each node $v \in G$, the MWC mechanism holds a contest for each subgraph $G_v = (V', E')$ rooted at v . If player v 's virtual credits $b_v \geq \eta \cdot (t_v)^2$, he will be allowed to enter into the contests. Otherwise, player v makes no contributions and receives 0 rewards. If v is allowed to enter into the contest, his probability of winning is determined by a contest success function: $prob(v) = \frac{(b_v)^\sigma}{\sum_{u \in V'} (b_u)^\sigma}$, where $0 < \sigma \leq 1$, and $b_v \geq \eta \cdot (t_v)^2$.

The parameter σ can be interpreted as the "noise" of a contest. It captures the marginal increase in the probability of winning caused by a higher effort and is crucial to the outcomes of the contest (Jia, Skaperdas, and Vaidya 2013). Contests with low σ can be regarded as poorly discriminating or "noisy" contests. That is, players with different efforts may have a similar level of chance to win. Contests with high σ can be regarded as highly discriminating in the sense that players with higher efforts have a greater chance to win. The contest success functions are imperfectly discriminating in the sense that with all of them, the prize at stake is awarded probabilistically to one of the players with higher

virtual credits leading to a higher probability of winning the prize.

In the contest success function, each player’s probability of success does not depend on his identity or the identities of his opponents, but just on the efforts (i.e., virtual credits) of the players. This property indicates that if in some cases two players have identical efforts then their probabilities of success must be equal and if all players were to exert identical efforts, then each one of them would have a probability of success equal to $1/n$. Besides, a player’s probability of success is independent of agents who have not exerted efforts.

The total rewards for diffusion contributions are $\pi_d = \phi \cdot \sum_{u \in V'} t_u \frac{deg_{G_v}^-(u)}{deg^-(u)}$, where $\phi > 0$, and $deg_{G_v}^-(u)$ denotes the number of direct predecessors of u in $G_v = (V', E')$. Here, $deg_{G_v}^-(u) \leq deg^-(u)$. The diffusion rewards of v are thus determined as:

$$\pi_d(v) = \frac{(b_v)^\sigma}{\sum_{u \in V'} (b_u)^\sigma} \cdot \phi \cdot \sum_{u \in V'} t_u \cdot \frac{deg_{G_v}^-(u)}{deg^-(u)}. \quad (3)$$

The MWC mechanism utilizes the probability of winning to determine the proportion of the total diffusion rewards for each player. This is because empirical evidence shows that proportional-prize contests exert higher aggregated efforts than winner-take-all contests (Cason, Masters, and Sheremeta 2010; Sheremeta 2011). Besides, proportional-prize contests also limit the degree of biases that discourage low-ability players without altering the performance of stronger players (Cason, Masters, and Sheremeta 2010; 2018).

The MWC Mechanism

The MWC mechanism $\pi = (\pi(v))_{v \in G}$ uses a post-price mechanism $\pi_t(v) = \mu \cdot t_v$ to reward a player that exerts t_v task efforts. The total reward for player v includes both the task rewards $\pi_t(v)$ and the diffusion rewards $\pi_d(v)$. Thus, we have the total reward for player v :

$$\pi(v) = \begin{cases} \mu \cdot t_v + \pi_d(v) & \text{if } b_v \geq \eta \cdot (t_v)^2 \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

where $\mu \in \mathbb{R}_{>0}$ is the reward parameter that characterizes to what extent the principal values agents’ efforts, and b_v , $\pi_d(v)$ are defined by Equations 2 and 3, respectively.

The MWC mechanism π (See Algorithm 1 for an implementation) is computationally efficient. The social network $G = (V, E)$ can be implemented using adjacency lists, where each node maintains a list of all its adjacent edges. It takes $O(|V|^2 + |V| \cdot |E|)$ time to compute the virtual credits (See Lines 1–6). Computing the total rewards (See Lines 7–9) also takes $O(|V|^2 + |V| \cdot |E|)$ time. Therefore, the time complexity for the MWC mechanism is $O(|V|^2 + |V| \cdot |E|)$.

The MWC mechanism achieves false-name-proofness by introducing an initial allowance $\eta_v \cdot (t_v)^2$ of virtual credits for each player v that has exerted non-zero task efforts. If the task efforts reduce, the allowance will decrease quadratically. False-name-attacks are not profitable because splitting task efforts into smaller pieces causes a larger degree of reduction than the degree of increase in virtual credits.

Algorithm 1 Multi-Winner Contests Mechanism

Input: G - a social network; S -seed nodes
Output: π -payment for each node $v \in G$
Initialize: $\pi(v) = 0$ for all $v \in G$
1: **for** each newly joined player $v \in G$ **and** $t_v > 0$ **do**
2: $b_v \leftarrow \eta \cdot (t_v)^2$
3: **for** each $v_i \in \kappa_v^-$ **and** $t_{v_i} > 0$ **do**
4: $b_{v_i} \leftarrow b_{v_i} + t_{v_i} \cdot t_v \cdot \sum_{p \in P_{v_i v}} w(p) \cdot \lambda^{|p|}$
5: **end for**
6: **end for**
7: **for** each v_j in G_v **do**
8: Compute the rewards $\pi(v_j)$ by Equation 4
9: **end for**
10: **return** π

This elegant design of the virtual credits enables the MWC mechanism to be manipulation-resistant to all the three types of false-name attacks. We postpone all the proofs to the Appendix.

Theorem 1 (False-Name-Proofness). *The MWC mechanism π is false-name-proof.*

The MWC mechanism employs a post-price mechanism π_t to determine the rewards for players’ task efforts. For the diffusion rewards, each player that has made successful referrals is guaranteed to receive a positive amount of rewards. Thus, it follows that the MWC mechanism satisfies the individual rationality property.

Theorem 2 (Individual Rationality). *The MWC mechanism π is individually rational.*

The MWC mechanism is budget-constrained because it allocates a fixed portion of rewards (i.e., $\mu \cdot \sum_{v \in G} t_v$) for the task efforts and a fixed number of rewards (i.e., $\phi \cdot \sum_{v \in G} t_v$) for diffusion rewards.

Theorem 3 (Budget Constraint). *The MWC mechanism π is budget-constrained.*

Under the MWC mechanism, players with direct referrals are given more virtual credits than if they have indirect referrals. This is achieved by introducing the discounting factor $\lambda < 1$. For the same level of task efforts, the virtual credits for diffusion rewards decrease as the distance of path between the referrer and the successor increases.

Theorem 4 (Monotonicity). *The MWC mechanism π is monotonic.*

Since the diffusion rewards are determined by a contest among players in a subgraph rooted at player v , it follows that the MWC mechanism is subgraph constrained.

Theorem 5 (Subgraph Constraint). *The MWC mechanism π satisfies the subgraph constraint.*

Experiments

Before describing the results, we introduce the experimental settings.

Experimental Settings

We used four publicly available datasets: Twitter (Hodas and Lerman 2014), Flickr (Cha, Mislove, and Gummadi 2009), Flixster (Goyal, Bonchi, and Lakshmanan 2011), and Digg (Hogg and Lerman 2012). These datasets included anonymized timestamps that could be used to estimate the influence probabilities needed for simulating the diffusion process of social networks (Goyal, Bonchi, and Lakshmanan 2010). For each dataset, we first estimated the influence diffusion probabilities for each node using the learning algorithms by Goyal, Bonchi, and Lakshmanan (2010) with the Bernoulli distribution under the static model. We then simulated the influence diffusion process with the general threshold model (Kempe, Kleinberg, and Tardos 2003) based on the estimated diffusion probabilities. After preprocessing, each dataset produced a largest weakly connected component. See Table 1 for dataset configuration.

Dataset	#Nodes	#Edges	#Seeds	M.D.	A.D.
Twitter	323,185	2,148,717	1,715	8,822	52
Flickr	145,305	2,149,882	768	6,731	34
Flixster	95,969	484,865	502	3,109	27
Digg	17,817	128,587	107	1,375	20

Table 1: Dataset configuration: M.D. – maximum degree, A.D. – average degree.

We modeled players’ abilities with the simulation method by Burnap et al. (2013). We considered four groups of players: *homogeneous (HO)*-players with similar levels of abilities; *heterogeneously low (HL)*-players with different level of abilities, and the average abilities are low; *heterogeneously high (HH)*-players with different level of abilities, and the average abilities are high; and *distinct (DI)*- a portion of players with low average abilities, the other portion with high average abilities. Each player’s ability was generated according to a Gaussian distribution with means that follow the probability density function (See Figure 3). Let ρ_v be player v ’s ability, then his cost coefficient $\delta_v = 1 - \rho_v$.

In our experiments, we set $\lambda = 0.5$ as it was standard in many geometric reward mechanisms. In practice, a stakeholder usually sets $\varphi \leq 1$ to make profits, but φ should be as close to 1 as possible to encourage players to participate. We let $\varphi = 1$. To encourage players to join, we set $\mu = 0.9$, and $\phi = \varphi - \mu = 0.1$. Note that $\eta \geq \lambda/2 = 0.25$, we set $\eta = 0.25$. For each group of players in each dataset, we varied the noise factors from 0 to 1 with an increment of 0.05. For each result (i.e., a data point) obtained, we ran the respective experiment 20 times. We ran all the experiments on the same 3.7GHz 6-core Linux machine with 32GB RAM.

Results

For each dataset, we compared the total aggregated task efforts by each group of players as the noise factor σ varied (See Figure 4). Figure 4a shows that the total contributions were quite low ($\leq 2,000$) when the noise factor was zero. The main reason is that players had no incentives to make

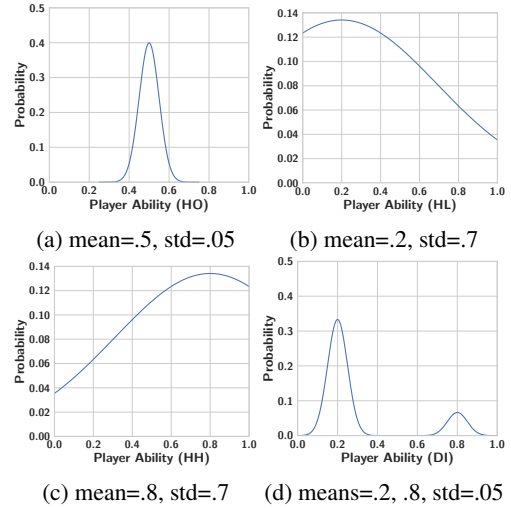


Figure 3: Probability density function of the means of players’ abilities by four different groups: HO, HL, HH and DI.

extra diffusion contributions (See Figure 5a) since the diffusion rewards were determined by a random lottery with no dependence upon the efforts of players when $\sigma = 0$.

As σ began to increase, the total contributions experienced a significant growth, which was largely due to the increase of players (See Figure 5a). The growth continued until it reached the peak where the total contributions began to shrink gradually (See Figure 4). The decline continued until it started to level off when the noise factor exceeded 0.8 (See Figure 4). An explanation for this trend is that when σ increased, the MWC mechanism values higher diffusion efforts more. When σ is low, the reward mechanism favors players with average abilities. These players usually take up a large portion of the population. When σ reached a point, the mechanism no longer leaned toward these players. It became more appealing to high-ability players than the low-ability or ordinary-ability players (See Figure 5b). However, all groups except the HH group included a small portion of high-ability players. That also explains why the HH group experienced a slighter decline than the other groups. Similar observations were found in social networks.

Players in the HH and HO groups typically performed better than the other two groups. As the noise factor σ increased, the HO groups experienced the highest degree of growth in total task efforts. Two factors contributed to this phenomenon. First, HO groups mainly consisted of players with normal abilities (See Figure 3a). Second, when σ increased to a point, the contests favored players with common abilities. Figure 4 shows that the total contributions of the DI groups were the most invariant to the noise factors. This is because the DI groups lacked a population mass of players with ordinary abilities (See Figure 3d). The observations applied to all the three datasets (See Figure 4). Despite that the optimal noise factors differed in populations and network structures, there were “sweet spots” for stakeholders to maximize the total efforts. In our experiments, all the optimal noise factors fell within the range between 0.45




Figure 4: A comparison of total contributions with different noise factors.

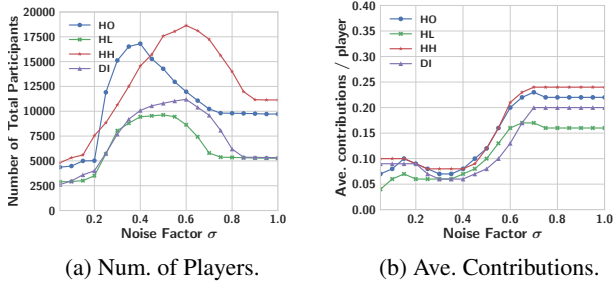


Figure 5: A comparison of number of participants and average contributions per player on Twitter.

and 0.65 (See Figure 4), which suggested that medium noise factors were typically superior than the others.

To study how false-name attacks affect players' rewards, we measured the total rewards earned by a player that had created different numbers of false identities. For each network, we considered a population consisting of the same percentage (25%) of players from each of the four groups: HO, HL, HH, and DI. We compared four noise factors: 0.4, 0.5, 0.6, and 0.7 with all the four networks.

Results show that the rewards earned by each player reduced significantly as the number of false identities increased from zero. The decline then slowed down until the rewards remained almost steady (See Figure 6). These trends were observed on all the groups of experiments that were given different noise factors and different network structures. An explanation is that when a player created a small number of false identities, his task efforts would be diluted to each false identity. This substantially reduced the player's virtual credits and his diffusion rewards. When the number of false identities continued to increase, there was little room to reduce. This consistent results further echoed our theoret-

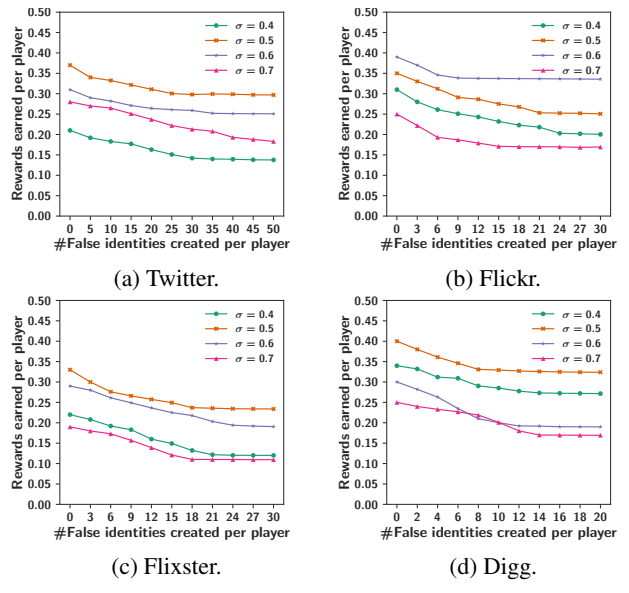


Figure 6: Rewards (normalized) obtained by a player that creates different number of false identities for groups with different noise factors.

ical analysis that the MWC mechanism is false-name-proof.

In summary, experimental results indicate that stakeholders can maximize the total task contributions by selecting appropriate (typically moderate) noise factors for the contests. It further demonstrates that players will not gain by creating false identities. The results also show that the MWC mechanism can scale to large social networks with hundreds of thousands of nodes.

Conclusion

In this paper, we have introduced a novel multi-winner contests mechanism for strategic diffusion in social networks. The mechanism is false-name-proof and individual rational for players with successful referrals. It is computationally efficient, budget-constrained, monotonic, and subgraph-constrained. Experiments on four real-world social networking datasets show that stakeholders can boost the performance of players by selecting proper noise factors. They further indicate that the MWC mechanism is both false-name-proof and scalable. Our work demonstrates the promising prospects of bringing contests to mechanism design.

Our work opens several exciting avenues for future research. In our model, agents do not discount the future and do not have uncertainties about the delivery of the rewards. In many real-world activities (e.g., crowdfunding, and investments), however, agents need to consider future risks when they make decisions. Another fertile area is to develop novel methods to automatically select optimal noise factors of MWC mechanisms for different concerns (e.g., profit maximization). We also find it very rewarding to integrate contests into the design of truthful mechanisms (Shen, Lopes, and Crandall 2016; Zhao et al. 2018).

Appendix

Proof of Theorem 1

Proof. By Definition 2, the task rewards for player v remain unchanged even if he creates replicas. To prove $\pi(v) \geq \sum_{r \in R} \pi(r)$, it suffices that if the v 's diffusion rewards are greater than or equal to the sum of diffusion reward of her false identities. That is, $\pi_d(v) \geq \sum_{r \in R} \pi_d(r)$. To prove this condition holds, we first show that the virtual credits earned by v are strictly greater than the sum of the virtual credits earned by the replicas R . That is, $b_v > \sum_{r \in R} b_r$. There are three scenarios according the types of false-name attacks: type 1 attacks, type 2 attacks, and hybrid attacks. We prove each case in turn:

- **Type 1 attacks:** By Equation 2, the virtual credits of original node v : $b_v = \eta \cdot (t_v)^2 + t_v \cdot \sum_{u \in \kappa_v^+} \sum_{p \in P_{vu}} t_u \cdot \omega(p) \cdot \lambda^{|p|}$. For simplicity, we let $b(v) = \eta \cdot (t_v)^2 + t_v \cdot C(v)$. Note that node v shares the same successors with the replica r_m . Therefore, we have the virtual credits for replica r_m : $b_{r_m} = \eta \cdot (t_{r_m})^2 + t_{r_m} \cdot \sum_{u \in \kappa_r^+} \sum_{p \in P_{ru}} t_u \cdot \omega(p) \cdot \lambda^{|p|} = \eta \cdot (t_{r_m})^2 + t_{r_m} \cdot \sum_{u \in \kappa_v^+} \sum_{p \in P_{vu}} t_u \cdot \omega(p) \cdot \lambda^{|p|} = \eta \cdot (t_{r_m})^2 + t_{r_m} \cdot C(v)$.

Similarly, we have r_{m-1} : $b_{r_{m-1}} = \eta \cdot (t_{r_{m-1}})^2 + t_{r_{m-1}} \cdot \sum_{u \in \kappa_{r_{m-1}}^+} \sum_{p \in P_{r_{m-1}u}} t_u \cdot \omega(p) \cdot \lambda^{|p|} = \eta \cdot (t_{r_{m-1}})^2 + t_{r_{m-1}} \cdot \sum_{u \in \kappa_{r_m}^+} \sum_{p \in P_{r_mu}} t_u \cdot \omega(p) \cdot \lambda^{|p|} + \lambda t_{r_{m-1}} \cdot t_{r_m} \cdot \eta \cdot (t_{r_{m-1}})^2 + \lambda t_{r_{m-1}} \cdot C(v) + \lambda \cdot t_{r_{m-1}} \cdot t_{r_m}$.

Likewise, we have: $b_{r_{m-2}} = \eta \cdot (t_{r_{m-2}})^2 + \lambda^2 t_{r_{m-2}} \cdot C(v) + \lambda^2 t_{r_{m-2}} \cdot t_{r_m} + \lambda t_{r_{m-1}} t_{r_{m-2}}, \dots, b_{v_1} = \eta \cdot (t_{r_1})^2 + \lambda^{m-1} t_{r_1} \cdot C(v) + \lambda t_{r_1} \cdot t_{r_2} + \lambda^2 \cdot t_{r_1} \cdot t_{r_3} + \dots + \lambda^{m-1} t_{r_1} \cdot t_{r_m}$.

Now we have the sum of the virtual credits for all replicas:

$$\begin{aligned} \sum_{r \in R} b_r &= b_{r_1} + \dots + b_{r_{m-1}} + b_{b_m} \\ &= \eta \cdot [(t_{r_1})^2 + \dots + (t_{r_{m-1}})^2 + (t_{r_m})^2] \\ &\quad [\lambda^{m-1} t_{r_1} + \lambda^{m-2} t_{r_2} + \dots + \lambda t_{r_{m-1}} + t_{r_m}] + t_{r_m} \cdot \\ &\quad C(v) + \lambda t_{r_1} \cdot t_{r_2} + \lambda^2 \cdot t_{r_1} \cdot t_{r_3} + \dots + \lambda t_{r_{m-1}} \cdot t_{r_m} \\ &< \eta \cdot [(t_{r_1})^2 + \dots + (t_{r_{m-1}})^2 + (t_{r_m})^2] + [t_{r_1} + \\ &\quad t_{r_2} + \dots + t_{r_m}] \cdot C(v) + t_{r_1} \cdot t_{r_2} + \dots + t_{r_{m-1}} \cdot t_{r_m} \\ &= \eta [t_{r_1} + t_{r_2} + \dots + t_{r_m}]^2 + t_v \cdot C(v) + \\ &\quad (\lambda - 2\eta) [t_{r_1} \cdot t_{r_2} + \dots + t_{r_{m-1}} \cdot t_{r_m}] \\ &\leq \eta (t_v)^2 + t_v \cdot C(v) \\ &= b(v). \end{aligned}$$

The first inequality holds since $0 < \lambda < 1$. The second inequality holds because $\lambda \geq \eta/2$. Thus, $b_v > \sum_{r \in R} b_r$.

- **Type 2 attacks:** Note that replicas r_1, \dots, r_m share the same successors with node v :

$$\begin{aligned} b_{r_j} &= \eta \cdot (t_{r_j})^2 + t_{r_j} \cdot \sum_{u \in \kappa_{r_j}^+} \sum_{p \in P_{r_j u}} t_u \cdot \omega(p) \cdot \lambda^{|u|} \\ &= \eta \cdot (t_{r_j})^2 + t_{r_j} \cdot \sum_{u \in \kappa_v^+} \sum_{p \in P_{vu}} t_u \cdot \omega(p) \cdot \lambda^{|p|} \\ &= \eta \cdot (t_{r_j})^2 + t_{r_j} \cdot C(v). \end{aligned}$$

Thus, we have:

$$\begin{aligned} \sum_{r \in R} b_r &= \eta \cdot [(t_{r_1})^2 + \dots + (t_{r_{m-1}})^2 + (t_{r_m})^2] + \\ &\quad [t_{r_1} + t_{r_2} + \dots + t_{r_m}] \cdot C(v) \\ &< \eta \cdot [t_{r_1} + t_{r_2} + \dots + t_{r_m}]^2 + t_v \cdot C(v) \\ &= b_v. \end{aligned}$$

The inequality is by Multinomial Theorem. Thus, $b_v > \sum_{r \in R} b_r$.

- **Hybrid attacks:** Since hybrid attacks are combinations of type 1 attacks, and type 2 attacks, it is trivial to know that under hybrid attacks $b_v > \sum_{r \in R} b_r$.

To prove $\pi_d(v) \geq \sum_{r \in R} \pi_d(r)$, it suffices if the following inequality holds:

$$\frac{(b_v)^\sigma}{\sum_{u \in V_{G_v}} (b_u)^\sigma} > \sum_{r \in R} \frac{(b_r)^\sigma}{\sum_{u \in V_{G_r}} (b_u)^\sigma}. \quad (5)$$

Note that $G_r = G_v \setminus \{v\} \cup R$. Let $x = (b_v)^\sigma$, $y = \sum_{u \in V_{G_v}} (b_u)^\sigma$, $z = \sum_{r \in R} (b_r)^\sigma$, and $q = \sum_{u \in V_{G_r}} (b_u)^\sigma$, we have $0 < \frac{x}{y} < 1$, $x > z > 0$, and $y > q > 0$. Equation 5 is equivalent to: $\frac{x}{y} > \frac{z}{q}$. Since $0 < \frac{x}{y} < 1$, we have:

$$\frac{z}{q} \cdot \frac{x}{y} = \frac{x - d}{y - d} \cdot \frac{x}{y} = \frac{(x - d)y}{x(y - d)} = \frac{1 - d/x}{1 - d/y} < 1. \quad (6)$$

Thus, $\frac{x}{y} > \frac{z}{q}$. Therefore, $\pi_d(v) \geq \sum_{r \in R} \pi_d(r)$. Now it follows that the reward mechanism π is false-name proof. \square

Proof of Theorem 2

Proof. By Equations 1 and 4, we have player v 's expected utility: $U(v) = \mu \cdot t_v + \pi_d(v) - \delta_v \cdot t_v = (\mu - \delta_v) \cdot t_v + \pi_d(v)$. Since the parameter μ is known to the player in advance, $\mu - \delta_v \geq 0$. Since $\pi_d(v) > 0$ if player has exerted diffusion contributions, we have $U(v) > 0$ for all players that have exerted both task efforts and diffusion contributions. \square

Proof of Theorem 3

Proof. By Equation 4, $\sum_{v \in G} \pi(v) = \sum_{v \in G} (\mu \cdot t_v + \pi_d(v))$. By Equation 3, $\sum_{v \in G} \pi_d(v) \leq \phi \sum_{v \in G} t_v$. Thus, we have $\sum_{v \in G} \pi(v) \leq (\mu + \phi) \sum_{v \in G} t_v$. Let $\vartheta = \mu + \phi$, we have $\sum_{v \in G} \pi(v) \leq \vartheta \sum_{v \in G} t_v$. \square

Proof of Theorem 4

Proof. Consider that the virtual credits of v_1 is $b_{v_1} = B$ when no successor is added. Now adding a direct successor v_i to v_2 , v_1 's virtual credits $b_{v_1'} = B + \sum_{p \in P_{v_1 v_i}} \lambda^{|p|} \cdot t_{v_1} \cdot t_{v_i} \cdot \omega(p)$. Similarly, if adding a direct successor of a successor of v_2 , we have v_1 's virtual credits: $b_{v_1''} = B + \sum_{p \in P_{v_1 v_j}} \lambda^{|p|} \cdot t_{v_1} \cdot t_{v_j} \cdot \omega(p)$. Note that $dist(v_1, v_i) < dist(v_1, v_j)$, and $t_{v_i} = t_{v_j}$, it suffices to show that $b_{v_1'} > b_{v_1''}$. Following the same method shown in Equation 6, we have the diffusion rewards: $\pi_d(v_1') \geq \pi_d(v_1'')$. \square

Proof of Theorem 5

Proof. According to Equation 3, player v virtual credits and consequentially the diffusion rewards are determined by a contest among players in the graph rooted at v , G_v . It follows that the MWC mechanism satisfies the subgraph constraint. \square

References

- Brill, M.; Conitzer, V.; Freeman, R.; and Shah, N. 2016. False-name-proof recommendations in social networks. In *AAMAS 2016*, 332–340.
- Burnap, A.; Ren, Y.; Papalambros, P. Y.; Gonzalez, R.; and Gerth, R. 2013. A simulation based estimation of crowd ability and its influence on crowdsourced evaluation of design concepts. In *the ASME 2013 International Design Engineering Technical Conferences*.
- Cason, T. N.; Masters, W. A.; and Sheremeta, R. M. 2010. Entry into winner-take-all and proportional-prize contests: An experimental study. *Journal of Public Economics* 94(9-10):604–611.
- Cason, T. N.; Masters, W.; and Sheremeta, R. 2018. Winner-take-all and proportional-prize contests: theory and experimental results. *Journal of Economic Behavior and Organization*.
- Cha, M.; Mislove, A.; and Gummadi, K. P. 2009. A measurement-driven analysis of information propagation in the flickr social network. In *WWW 2009*, 721–730.
- Chaffey, D. 2016. Global social media research summary 2016. *Smart Insights: Social Media Marketing*.
- Conitzer, V.; Immorlica, N.; Letchford, J.; Munagala, K.; and Wagman, L. 2010. False-name-proofness in social networks. In *International Workshop on Internet and Network Economics*, 209–221. Springer.
- Dow, S.; Kulkarni, A.; Klemmer, S.; and Hartmann, B. 2012. Shepherding the crowd yields better work. In *CSCW 2012*, 1013–1022.
- Drucker, F. A., and Fleischer, L. K. 2012. Simpler sybil-proof mechanisms for multi-level marketing. In *EC 2012*, 441–458.
- Emek, Y.; Karidi, R.; Tennenholtz, M.; and Zohar, A. 2011. Mechanisms for multi-level marketing. In *EC 2011*, 209–218. ACM.
- Ferrara, E.; Varol, O.; Davis, C.; Menczer, F.; and Flammini, A. 2016. The rise of social bots. *Communications of the ACM* 59(7):96–104.
- Galeotti, A., and Goyal, S. 2009. Influencing the influencers: a theory of strategic diffusion. *The RAND Journal of Economics* 40(3):509–532.
- Gao, H.; Liu, C. H.; Wang, W.; Zhao, J. R.; Song, Z.; Su, X.; Crowcroft, J.; and Leung, K. K. 2015. A survey of incentive mechanisms for participatory sensing. *IEEE Communications Surveys and Tutorials* 17(2):918–943.
- Goyal, A.; Bonchi, F.; and Lakshmanan, L. V. 2010. Learning influence probabilities in social networks. In *WSDM 2010*, 241–250.
- Goyal, A.; Bonchi, F.; and Lakshmanan, L. V. 2011. A data-based approach to social influence maximization. *the VLDB Endowment* 5(1):73–84.
- Hodas, N. O., and Lerman, K. 2014. The simple rules of social contagion. *Scientific Reports* 4:4343.
- Hogg, T., and Lerman, K. 2012. Social dynamics of digg. *EPJ Data Science* 1(1):5.
- Ipeiritos, P. G. 2010. Analyzing the amazon mechanical turk marketplace. *XRDS: Crossroads, The ACM Magazine for Students* 17(2):16–21.
- Jackson, M. O., and Yariv, L. 2011. Diffusion, strategic interaction, and social structure. In *Handbook of Social Economics*, volume 1. Elsevier. 645–678.
- Jia, H.; Skaperdas, S.; and Vaidya, S. 2013. Contest functions: Theoretical foundations and issues in estimation. *Int. J. Ind. Organ.* 31(3):211–222.
- Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *KDD 2003*, 137–146.
- Leskovec, J.; Adamic, L. A.; and Huberman, B. A. 2007. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)* 1(1):5.
- Lorenz, J.; Rauhut, H.; Schweitzer, F.; and Helbing, D. 2011. How social influence can undermine the wisdom of crowd effect. *Proc. Natl. Acad. Sci. USA* 108(22):9020–9025.
- Naroditskiy, V.; Jennings, N. R.; Van Hentenryck, P.; and Cebrian, M. 2014a. Crowdsourcing contest dilemma. *J R Soc Interface* 11(99):20140532.
- Naroditskiy, V.; Stein, S.; Tonin, M.; Tran-Thanh, L.; Vlassopoulos, M.; and Jennings, N. R. 2014b. Referral incentives in crowd-funding. In *HCOMP 2014*, 171–183.
- Pickard, G.; Pan, W.; Rahwan, I.; Cebrian, M.; Crane, R.; Madan, A.; and Pentland, A. 2011. Time-critical social mobilization. *Science* 334(6055):509–512.
- Rahwan, I.; Dsouza, S.; Rutherford, A.; Naroditskiy, V.; McInerney, J.; Venanzi, M.; Jennings, N. R.; and Cebrian, M. 2013. Global manhunt pushes the limits of social mobilization. *Computer* 46(4):68–75.
- Shen, W.; Crandall, J. W.; Yan, K.; and Lopes, C. V. 2018. Information design in crowdfunding under thresholding policies. In *AAMAS 2018*, 632–640.
- Shen, W.; Lopes, C. V.; and Crandall, J. W. 2016. An online mechanism for ridesharing in autonomous mobility-on-demand systems. In *IJCAI 2016*, 475–481.
- Sheremeta, R. M. 2011. Contest design: An experimental investigation. *Economic Inquiry* 49(2):573–590.
- Skaperdas, S. 1996. Contest success functions. *Economic Theory* 7(2):283–290.
- Tang, J. C.; Cebrian, M.; Giacobbe, N. A.; Kim, H.-W.; Kim, T.; and Wickert, D. B. 2011. Reflecting on the darpa red balloon challenge. *Communications of the ACM* 54(4):78–85.
- Todo, T.; Iwasaki, A.; and Yokoo, M. 2011. False-name-proof mechanism design without money. In *AAMAS 2011*, 651–658.
- Zhao, D.; Li, B.; Xu, J.; Hao, D.; and Jennings, N. R. 2018. Selling multiple items via social networks. In *AAMAS 2018*, 68–76.