

Game-theoretic Models of Information Overload in Social Networks^{*}

Christian Borgs^{**}, Jennifer Chayes^{**}, Brian Karrer^{***}, Brendan Meeder[†], R. Ravi[†],
Ray Reagans[‡], and Amin Sayedi[†]

Microsoft Research New England, Cambridge, MA

Abstract. We study the effect of information overload on user engagement in an asymmetric social network like Twitter. We introduce simple game-theoretic models that capture rate competition between celebrities producing updates in such networks where users non-strategically choose a subset of celebrities to follow based on the utility derived from high quality updates as well as disutility derived from having to wade through too many updates. Our two variants model the two behaviors of users dropping some potential connections (followership model) or leaving the network altogether (engagement model). We show that under a simple formulation of celebrity rate competition, there is no pure strategy Nash equilibrium under the first model. We then identify special cases in both models when pure rate equilibria exist for the celebrities: For the followership model, we show existence of a pure rate equilibrium when there is a global ranking of the celebrities in terms of the quality of their updates to users. This result also generalizes to the case when there is a partial order consistent with all the linear orders of the celebrities based on their qualities to the users. Furthermore, these equilibria can be computed in polynomial time. For the engagement model, pure rate equilibria exist when all users are interested in the same number of celebrities, or when they are interested in at most two. Finally, we also give a finite though inefficient procedure to determine if pure equilibria exist in the general case of the followership model.

1 Introduction

Social networking sites such as Facebook and Twitter allow users to sign up and keep in touch with others by friending them and receiving updates from their friends at their convenience by logging into web pages that provide the current feed of most recent updates; various mobile applications also port this convenience to smart phones. The two conceptually different modes of establishing this online relationship are symmetric or asymmetric, requiring consent from both sides or only from the following side to initiate and maintain the tie. These two modes are reflected in two major social network platforms of today: Facebook and Twitter. There are also other major differences between these two platforms in that Facebook is a full-fledged sharing site that allows users to share photos and longer updates, while Twitter is a micro-blogging site restricting each update to a text message of up to 140 characters. These design choices drive the differences in user enrollment and engagement of users between these two platforms (see, e.g., the infographic in [1]) with the convenience of small tweets and asymmetric nature of Twitter giving very high rates of adoption, but the more substantial updates and symmetric consent maintaining more users in Facebook.

Despite these differences, one feature is common to the basic versions of both home pages: In both platforms, users see one linear update feed reflecting a chronologically sorted order of posts from

^{*} Work done when all authors were visiting Microsoft Research New England in the summer of 2010.

^{**} Microsoft Research New England

^{***} University of Michigan, Ann Arbor

[†] Carnegie Mellon University

[‡] MIT Sloan School of Management

friends whenever she logs into each of these sites¹. Given that the foremost advantage of these sites is the convenience of getting a quick asynchronous update, the balance of information from various friends in the feed becomes important in determining the value the user will gain from the update. However, the mix of the current feed, such as the proportion from various friends is determined by the activity level of these friends. This leads to a fundamental shortcoming of the online medium in replicating real-world ties: even though the time and frequency of updating oneself is in one’s own control, how much one hears from one particular friend is not ². We model the implication of two natural control strategies that users adopt in response (i.e., unfollowing and disengagement) in this paper.

The lack of control of regulating the rate of flow of information in a particular link to a tie is part of the broader difficulty of online social networks replicating the off-line relationships: in the latter, even with asynchronous forms of communication such as postal and electronic mail, as well as in the more familiar forms of synchronized communication such as telephone or face-to-face conversations, the amount of time and hence, information communicated can be regulated at will by both parties. Even in the asymmetric follow cases of this such as attending a lecture in a conference, the user usually has the cognitive ability to switch off from the information presented to her and focus on other activities. However, the current design of online information update systems have reduced the number of degrees of freedom of a user in sending information out to her friends from the number of friends down to essentially one³. Our study can be seen as a first step in examining the effects of this key design decision.

1.1 Modeling the resulting networks

In this paper, we propose two models for explaining the set of ties that are realized and stabilized over time in an asymmetric follow network such as Twitter. This allows us to restrict our attention solely to the update rate feature while not worrying about symmetric consent; furthermore, given that Facebook has already implemented filters for enhancing the update experience while Twitter has not yet done so, our models of information overload focus on the latter where they are more relevant. Both of our models assume that the rate of sending updates to followers is the key decision variable of an agent in the network to help maintain a good follower circle. Also, both models assume that updates from friends are useful to the user but an excessive rate of updates have diminishing returns and eventually have negative returns of value in terms of flooding out the stream of useful updates from other friends.

The first set of “followership” models assume that users in the network will stay in the network but will unfollow agents from whom the updates are too frequent. The second “engagement” model assumes that users are frustrated by the high update rate of their followees and leave the network altogether with probability proportional to the rate of annoyance from excessively updating friends. Both sets of models are static in that they assume a fixed set of agents; moreover, our models explicitly partition users as producers or consumers of information but not both — this is not accurate in that agents in social networks serve both functions but it is a useful enough abstraction for our purpose. Moreover, given that over 80% of the tweets come from under 20% of the users, the

¹ There are variants to this common feature as well: Facebook learns and filters the feed to show what it believes are the most interesting recent updates, and various Twitter clients that allow access to the tweets allow other view options, but we will stick to the basic unfiltered version here for our models.

² This also speaks to an important feature that is lacking in such sites and might greatly impact the ability of the user to tailor the feeds and hence stay engaged in the platform — this idea has already been formulated as information overload being a “filter failure problem” by Clay Shirkey, see [6].

³ There are exceptions to these too, of course, such as using lists in Twitter or groups in Facebook, but again our discussion applies to the plain version of these platforms that new users are exposed to.

bipartite model of agents separated loosely into “celebrities” that produce information and “users” that simply consume it is not a major deviation from reality.⁴

The information contained in specific posts is not included in our models although it can vary significantly (from being updates of key news and events, to information nuggets about specific topic areas or simple grooming updates). Instead, each user associates a value to updates from a celebrity that reflect how interested they are in hearing from that celebrity. This value reflects the average utility the user gets from reading an update from that celebrity and in this way, we avoid the specific content of posts.

In both sets of models, our goal is to build simple models of rate competition between celebrities and characterize the ties that remain in the network at equilibrium, and also to infer the “optimal” rates resulting from the model for the celebrities. This can shed light on the disparities between offline and online networks based on the above discussion, as well as on how far away from optimality current celebrity tweet rates can be. While there is anecdotal discussion in the media, especially pertaining to social media marketing, about producers of information (such as corporate Twitter accounts) being cognizant of the rate and quality of their updates, this strategic aspect of information producer behavior has not been studied carefully before our work.

1.2 Our models

Our basic models assume a complete bipartite graph on two disjoint sets of nodes: the producers or celebrities (denoted C) and the consumers or simply users or followers (denoted F). Every edge between celebrity i and user j comes with a non-negative quality score q_{ij} that determines the proportional utility that j derives from following celebrity i . If an edge has quality score zero, we will commonly treat that edge as non-existent. An edge between celebrity i and user j in this bipartite graph represents the *possibility* of j following i . In the game-theoretic models, the players are the celebrities that compete by setting nonnegative rates r_i , and deriving utility proportional to their influence: namely, celebrity i updating at rate r_i has payoff equal to r_i times the number of followers he has at this rate (which will depend on the rates of other celebrities as dictated by the particular model). The main difference between the two models is the behavior of the users who are simple payoff maximizers.

Followership Games: For the first model that we term “Followership Games”, the user’s payoff is the sum of quality-weighted update rates of the celebrities he follows minus $\lambda \geq 0$ times a superlinear function of the *total* rate of all celebrities he follows (for some parameter input λ) — we will henceforth use the quadratic function of the total rate but other choices such as the exponential or some polynomial can also be adopted. Higher λ ’s imply higher disutility from increasing total rate of updates of followed celebrities, and the choice of superlinear function also has a similar effect as the function magnifies the amplitude of the total updates the user receives. One can also use a sublinear function of the total rate of updates but unless λ is very large, this will result in all competing celebrities to update at unbounded rates in the resulting competition. Furthermore, using a linear rather than quadratic function for the disutility term leads to a decomposition of the user’s utility into one term per celebrity: when viewed from the celebrity’s side, this competition problem becomes trivial — every celebrity i faces a certain net linear response from each user j (with parameters varying based on the quality q_{ij} the user feels from this celebrity and λ) and has to choose an optimal rate of update r_i that maximizes her utility. This problem has no influence from the rates of other celebrities and hence does not involve any competition. For these reasons, we only focus on the superlinear, and in particular the quadratic disutility case.

Engagement Games: In the second model, the users are assumed to get disengaged and leave the network altogether if they receive too many updates from the celebrities. In terms of the bipartite graph, the user is assumed to be currently following all users for whom they have non-zero quality,

⁴ Applying our followership model to a non-bipartite network can be done by duplicating each individual into a celebrity and a user.

and for this paper, we assume all quality values are one in this model. The result of the Engagement Game determines whether the user continues to engage with all of the celebrities (for which their quality is non-zero) by following them, or they disengage and stop following all of them.

1.3 Our Results

Our main result is negative showing that even the simple followership model of rate competition does not admit pure Nash equilibria in general (Example 1). This comes from the fundamentally discrete and hence discontinuous nature of the utility that producers derive as they increase their rates: at certain breakpoints they start to lose followers for whom they exceed their “tolerance” limit for following them, and hence there are kinks in the producer utility function.

While we do not present necessary conditions for existence of pure equilibria, we describe several special cases with sufficient conditions: For the Followership model, one is when there is a global ranking of celebrities such that the quality ordering for any user follows the same order among the celebrities (Proposition 3); a slight generalization of this holds when a dependency graph among the celebrities (that reflects quality dominance for some user) is acyclic (Proposition 4). For the Engagement model, one case where pure rate equilibria exist is when every user has unit quality for exactly k celebrities and no quality from others (Proposition 5, and another is when every user has unit quality for at most two celebrities which is the case of users with sparse interests among the celebrities (Proposition 6).

In the Followership model, we use simple first order conditions to arrive at a finite set of linear conditions based on matchings to determine if there is a pure equilibrium (Section 2.3). These matching conditions characterize the set of all possible rates that are candidates for equilibrium. These conditions produce candidate rates to check for equilibrium. While not an efficient characterization, they do provide finite necessary characterization of pure strategy rate equilibria for the Followership game. We propose a similar set of linear conditions for the Engagement game based on sets of possibly engaged users, but are unable to verify that they characterize all pure equilibria.

1.4 Related Work

The concept of information overload has a long history in the management literature. Roughly speaking, an individual’s ability to process information for a task is supposed to be an inverted U-shape with respect to information quantity. Increasing information at first increases an individual’s capability but eventually additional information becomes unhelpful and information-processing ability declines. This general phenomena is seen in a wide variety of disciplines. [2]

Empirical studies have attempted to identify a “core group” of links in networks like Twitter that are active and matter [3] in capturing the real underlying activity in the network. Our work differs from these streams of work in that we focus on models for the emergence of such a core network in Twitter.

Many papers have addressed network formation games in other contexts [4], but have modeled the utility of a node from connecting to a neighbor in more coarse terms, without taking into account the temporal details of the interaction such as the rate of information flow that we attempt to model in our investigation.

Our models are related to problems of pricing between substitutable products by a monopolist: In this analogy, the updates represent quantities of products, celebrities represent substitutable products, followers represent consumers and the qualities of an update from a celebrity to a user resemble the utility multiplier for the product for the consumer. In this context there is much work on competitive pricing across products and how this results in over or under-entry of products into this market starting from the work of Sattinger [5]. However, a key difference in our models is the additional disutility effect which results from consuming too much.

1.5 Empirical Evidence

To support our work empirically we performed a large-scale analysis of over ten-million Twitter users. We collected the profile information of users who have unique user IDs between 20,000,000 and 30,000,000; these users created their accounts between February 3, 2009 and April 9, 2009. Within this set of users we focused on those who had between 1,000 and 5,000 followers on August 28, 2010. In our data set there are over 65,654 users meeting this condition, and we call such accounts ‘micro-celebrities.’ These users have possessed accounts for a long enough duration that we believe they are in some sort of equilibrium; they understand how Twitter works, have built up a substantial following, and have developed their own usage patterns. We recollected profile information and a list of followers for each micro-celebrity several times over the duration of two weeks.

We partitioned the users into buckets based on the number of tweets they created between August 28 and September 13, 2010. The buckets we used are exponentially sized: $B_0 = \{0\}$, $B_1 = \{1\}$, $B_2 = \{2, 3\}$, \dots , $B_{13} = [2^{12}, 2^{13} - 1]$. For each bucket, we looked at how many accounts unfollowed each user in the bucket and computed the mean, median, 25th, and 75th percentile of these unfollow counts. We present these numbers in the plot below; it is clear that as a micro-celebrity tweets at a higher rate, the number of accounts unfollowing that user increases.




Fig. 1. Number of unfollow events against the number of tweets authored in the observation period.

In the next two sections, we review our two classes of Followership and Engagement models and then conclude with some directions of current and future work.

2 A Model of Friendship Selection: the Followership Game

Assume that there are n celebrities and m users. Consider some celebrity i . Suppose that the rate of every other celebrity is fixed. We assume that the utility of a user for a celebrity as the celebrity changes her tweet-rate has an inverse U-shape as suggested by the information overload literature [2]; see Figure 2. The parameters of the curve, e.g., where it hit the axis, and its rate of change of slope, however depend on the user model and the tweet-rate of the other celebrities. Given these assumptions, it is clear that a user will not follow a celebrity whom he has negative utility for. Therefore, if an individual celebrity keeps increasing her tweet-rate, her followers will stop following her one by one.

If we assume that the utility of a celebrity is her *influence*, i.e., her tweet-rate times the number of followers she has, her utility-rate curve would look like Figure 3 with discontinuities at the points where her followers drop her. The curve has slope k in the first piece, where $k \leq m$ is the number of users who follow the celebrity if she has very low tweet-rate. The second piece has slope $k - 1$ and so on, until the last piece which has slope 1. If a celebrity wants to maximize her utility given the tweet-rate of the other celebrities, she has to find the peak of this curve and update at this rate.




Fig. 2. Utility of User for a Specific Celebrity




Fig. 3. Utility of a Celebrity

2.1 A Specific Utility Model

We assume a complete weighted bipartite graph between celebrities and users. For celebrity i and user j , the weight q_{ij} associated with edge (i, j) indicates the quality of user j for celebrity i . The higher q_{ij} is, the more utility user j gets from tweets of celebrity i . We assume that all q_{ij} 's are known to the celebrities and the users. As stated before, we define the utility of celebrity i to be $U_i = |F_i|r_i$ where F_i is the set of users who follow her, and r_i is her tweet-rate.

For the utility function of a user, we want some function that captures the property described in Figure 2. Moreover, we need the function to capture the competition between the celebrities. For example, suppose that user a really likes celebrity x but is also following celebrity y . If celebrity x increases her tweet-rate, since the rate of information intake for user a increases, he gets *overloaded*. As a result, he would consider dropping one of the celebrities that he is following. Since he likes celebrity x more than celebrity y chances are that celebrity y will be dropped. Let C_j be the set of celebrities that user j follows; we define the utility of user j to be

$$U_j = \sum_{i \in C_j} r_i q_{ij} - \lambda \left(\sum_{i \in C_j} r_i \right)^2.$$

The first term represents how much the user is benefiting from consuming tweets while the second term represents the *information overload* concept. By scaling the qualities without loss of generality we assume $\lambda = 1/2$.

Users constantly optimize their utility function, i.e., choose the set of celebrities that they want to follow. At the same time, the celebrities are adjusting their rates to maximize their utility. Fix user j and let x_i be the indicator variable for following celebrity i . User j wants to maximize $U_j = \sum_i x_i r_i q_{i,j} - \lambda (\sum_i x_i r_i)^2$ subject to $x_i \in \{0, 1\}$. This is a non-linear integer program and is generally hard to solve. To simplify, let's first consider the fractional version of the problem where $x_i \in [0, 1]$, i.e., the user can follow celebrities fractionally.

Fractional Following If celebrity i is followed fractionally ($0 < x_i < 1$) by user j , we must have $\frac{\partial U_j}{\partial x_i} = 0$. The expression simplifies to $q_{ij} = 2\lambda \sum_k x_k r_k = \sum_k x_k r_k$. For any celebrity l where $q_{lj} > q_{ij}$ we must have $x_l = 1$, and for any celebrity l' where $q_{l'j} < q_{ij}$ we must have $x_{l'} = 0$. The graphical representation of the optimum fractional solution is given in Figure 4. Assume that $q_1 \geq \dots \geq q_n$ is the sorted sequence of q_{1j}, \dots, q_{nj} . Each horizontal segment corresponds to a celebrity; the length of the segment is the tweet-rate of the celebrity and its height is the quality of the user for this celebrity. There is a dashed line with slope 1 that goes through origin. The user will follow the celebrities to the left of the dashed line. The celebrity who intersects the dashed line is followed fractionally.




Fig. 4. User's Strategy

The utility of celebrity i in the fractional setting should be adjusted to $U_i = \sum_j x_{ij} r_i$ where x_{ij} indicates the fraction of celebrity i that user j is following. The following proposition helps us to characterize the pure Nash equilibria of this Followership Game in the fractional setting.

Proposition 1. *Consider celebrity i and fix the tweet-rate of the other celebrities. If i increases her tweet-rate, her utility U_i will not decrease.*

Proof. Consider an arbitrary user j . If celebrity i changes her rate from r_i to αr_i ($\alpha > 1$), variable x_{ij} in the optimal solution of user j will change to a value of at least $\geq x_{ij}/\alpha$. Therefore, the expression $U_i = \sum_j x_{ij} r_i$ will not decrease. Furthermore, if for some user j and celebrity l we have $x_{lj} > 0$, and $q_{lj} > q_{ij}$, increasing r_i strictly increases U_i .

Proposition 1 suggests that in equilibrium, the celebrities are tweeting at a very high rate and each user follows only one (maybe even fractionally) celebrity. Since this equilibria is unrealistic, we consider alternative “integral” user behavior in the Followership Game.

2.2 Greedy Users

One way of modifying the fractional solution obtained in section 2.1 is to drop the only (if any) celebrity who is followed fractionally. More precisely, if the optimal fractional solution for user j is to follow celebrity i fractionally (we know that at most one such celebrity exists for each user), now we assume that he does not follow her anymore; i.e., we set all the variables $x_{ij} < 1$ to 0. This is equivalent to say that the user only follows those celebrities that completely lie to the left of the dashed line in Figure 4. The following is an equivalent definition for user's behavior in this model.

Definition 1. Consider user j and let $q_1 \geq \dots \geq q_n$ be the sorted order of q_{1j}, \dots, q_{nj} . Let k be the largest index such that $\sum_{i=1}^k r_i \leq q_k$. Under the greedy users model, user j follows the k celebrities for who he has highest quality and no one else.

Next, we show that pure strategy Nash equilibrium does not necessarily exist in this model.

Example 1. There are 3 celebrities x, y and z , and $3(1 + N)$ users. Three special users are labeled a, b and c . See Figure 1. Qualities for these users are as follows: $q_{ax} = k + 1, q_{az} = k, q_{bx} = k, q_{by} = k + 1, q_{cy} = k$, and $q_{cz} = k + 1$. The other $3N$ users are partitioned into three equal-size groups G_x, G_y and G_z . The N users in group G_i ($i \in \{x, y, z\}$) have quality 1 for celebrity i and quality 0 for other celebrities. All other qualities are 0. Edges with zero quality are not shown.




Fig. 5. Non-existence of Pure Nash Equilibrium

Proposition 2. If $2k - 2 > N > k + 1$, the graph depicted in Example 1 does not have a pure strategy Nash equilibrium.

Proof. Assume for sake of contradiction that a pure strategy Nash equilibrium exists. First note that if a celebrity i sets her tweet-rate to 1, she can have utility at least $N + 1$ (N from G_i , and 1 from the user who has quality $k + 1$ for her). Therefore, since a celebrity can have at most one follower if her rate is strictly above k , any rate $> k$ is strictly dominated. Also observe that with any rate > 1 , the celebrity can have at most two followers; therefore, any rate $1 < r < (N + 1)/2$ is also dominated. As a result, the tweet-rate of each celebrity, if a Nash equilibrium exists, should be either in the interval $(0, 1]$ or in the interval $[(N + 1)/2, k]$.

In a Nash equilibrium, either zero, one, two, or all of the celebrities are tweeting at rate $\in (0, 1]$. We prove case-by-case that none of these possibilities can work to conclude that no equilibrium exists. If all three of them are tweeting at rate $\in (0, 1]$, it is clearly beneficial for celebrity x (or y , or z by symmetry) to deviate to tweet-rate k because her utility would change from $N + 2$ to $2k$.

If at least two celebrities are tweeting at rate higher than 1, we show that one of them would benefit from changing her rate to 1. Because of the symmetric structure of the graph, assume without loss of generality that celebrities x and y are tweeting at rate $\geq (N + 1)/2 \geq k/2 + 1$. Since sum of their rates is larger than k , user b will only follow y which makes the utility of celebrity x less than

or equal to k . Therefore, she would benefit from deviating to tweet-rate 1 which makes her utility $N + 2$.

Finally, consider the case where only one celebrity is tweeting at rate $\in [(N + 1)/2, k]$; because of the symmetric structure of the graph without loss of generality assume that this celebrity is x . If y changes her tweet-rate to k , users b and c will both follow her. Therefore, her utility will increase to $2k$ which contradicts the equilibrium assumption. This completes the proof that no pure strategy Nash equilibrium exists for this example.

Although this example proves that a pure strategy Nash equilibrium does not always exist in the greedy version of the Followership Game, there are special cases for which we can prove the existence.

Pure Rate Equilibria Exist Under Global Ranking. Suppose that there is a global ranking of the celebrities such that every user prefers higher ranked celebrities to lower ranked celebrities. Without loss of generality assume that the ordering is $1, \dots, n$. More precisely, the global ranking assumption means that for every user j we have $q_{1j} \geq \dots \geq q_{nj}$.

Proposition 3. *A pure strategy Nash equilibrium exists for the followership game when there is a global ranking for the celebrities.*

Proof. We construct an equilibrium iteratively. We calculate the equilibrium rate r_i^* of celebrity i in step i . In step i , set the rates of celebrities $1, \dots, i - 1$ to their equilibrium rate which has been calculated in the previous steps; set the rate of celebrities $i + 1, \dots, n$ to 0. We can calculate the optimum rate \hat{r}_i of celebrity i when the rates of the others are given (it follows from computing the utility-rate curve of Figure 3). We simply set r_i^* to \hat{r}_i .

Note that if some celebrity $i' > i$ changes her rate, given that the rate of everyone else is fixed, the utility of celebrity i will not be affected. Therefore, if utility of celebrity i is optimized at r_i^* when $r_{i+1} = \dots = r_n = 0$, it is still optimized when r_{i+1}, \dots, r_n have any other value. Hence, we can argue that if celebrities $1, \dots, i - 1$ do not benefit from deviating from their rate, celebrity i would not benefit either. Using induction, we conclude that the rates r_i^* form a pure strategy Nash equilibrium of the game.

This global ranking is a special case of a more general result.

Proposition 4. *For a given followership game instance, create a dependency graph on the celebrity nodes as follows: let x have a directed edge to y if the non-zero quality of x is greater than the non-zero quality of y for some user u . If the resulting dependency graph is a directed acyclic graph, a pure strategy Nash equilibrium exists for the rate competition game.*

Proof. Note that an edge from x to y means therefore that y takes into account x 's rate when myopically maximizing their utility. Since the dependency graph is a directed acyclic graph, the nodes can be topologically ordered and then the above induction argument will also apply to that ordering.

2.3 Matchings characterize pure rate equilibria

For this part, assume that all quality values for a user are distinct and hence there is a strict quality ordering of the celebrities for a user. We show that we can enumerate over a finite set to check if there are any pure Nash equilibria for the rates.

Note that from our description of our response function, for every celebrity whose rate is at equilibrium, there are two possibilities. Either there exists a user for who the celebrity is critically tweeting: i.e., an increase in her rate implies this user will drop her, or the celebrity has rate zero. These saturated users are distinct for distinct celebrities in the greedy users model that we are

working with when the celebrity qualities from a user are distinct. Thus, there is a simple (albeit slow) procedure for finding a pure Nash equilibria of rates if one exists: First, find a matching from all subsets of celebrities (who are updating at a nonzero rate in a possible equilibrium) to users that are critical for them; Then, solve for the first order conditions being tight for this set of users with the matched celebrity being the last one followed by this user. If the resulting system is nonsingular, these rates give a possible Nash Equilibrium. We then check whether it satisfies the other equilibrium conditions (nonnegative rates and non-profitable myopic deviations) to confirm whether the candidate solution is indeed correct.

3 A Model of User Engagement: the Engagement Game

Suppose that user j follows the set of celebrities C_j and celebrity i is followed by the set of users F_i . Instead of considering strict quality values, in the Engagement Game we merely assume that a user j is currently following all celebrities for which they have non-zero quality and not following any others. Celebrity i then decides about her tweet-rate r_i , and user j decides whether to remain active on Twitter or to *disengage*. So unlike the Followership Game, where follow relationships resulted from the game, here follow relationships are fixed and users then decide to leave the network, rather than break ties. One way of interpreting this is that users joining Twitter follow some subset of celebrities based on their interests. After using the service for some time they might find that the combined rate of information from these celebrities is too much and decide to delete their account.

Let $F_i^* \subseteq F_i$ denote the set of engaged followers of celebrity i . The utility of celebrity i is $U_i = r_i |F_i^*|$, where the set F_i^* will be a random variable whose probability depends on celebrity rates. So we assume that celebrities maximize their expected utility. Tweeting too much increases r_i but at the same time decreases the expected size of F_i^* . The celebrities are competing for user's attention but in order to keep the users engaged, they have to avoid high tweet-rates.

We suppose that there is a function $S : [0, \infty) \rightarrow [0, 1]$ that maps the total rate of all celebrities followed by a user into a probability that the user stays in the social network and this function is the same for all users. This function is intended to capture the frustration of a typical user from receiving too many updates, and the function outcome as a probability adds heterogeneity among the individual users though they all have the same function. Thus, user j stays in the network with probability $S(\sum_{i \in C_j} r_i)$ and the (expected) utility of celebrity i is

$$U_i(\mathbf{r}) = r_i \left(\sum_{j \in F_i} S \left(\sum_{i' \in C_j} r_{i'} \right) \right)$$

Certain choices of S admit a closed form solution for the celebrity strategies. For example, if $S(r) = \exp(-r)$, then the rate of each celebrity becomes decoupled:

$$U_i(\mathbf{r}) = r_i \left(\sum_{j \in F_i} \exp \left(- \sum_{i' \in C_j} r_{i'} \right) \right) = r_i \sum_{j \in F_i} \exp(-r_i) \exp \left(- \sum_{i' \in C_j, i' \neq i} r_{i'} \right) = r_i \exp(-r_i) \sum_{j \in F_i} c_j$$

where $c_j = \exp(-\sum_{i' \in C_j, i' \neq i} r_{i'})$ is constant w.r.t. r_i . An easy calculation shows that the optimal solution is $r_i = 1$. Notice that even if any of the other $r_{i'}$ change, the optimal solution for producer i is still $r_i = 1$. Thus, for this particular S function all the producers tweet at rate 1. Instead of this factorized form, we assume a non-trivial function that causes interaction between producers.

Suppose that the function is $S(r) = \max(0, 1 - r)$. Then the probability that user j engages is $\max(0, 1 - \sum_{i \in C_j} r_i)$. In general, pure Nash equilibria for this game are characterized by a set of rates along with a subset F' of users j that have nonnegative values of $1 - \sum_{i \in C_j} r_i$. The pure equilibrium rate for a celebrity i is a solution to the constrained optimization problems of maximizing the celebrity's utility $r_i \cdot S(r_i)$ subject to the constraints that for every j in the subset F' the probability

expression $1 - \sum_{i \in C_j} r_i$ of staying engaged is nonnegative, while it is nonpositive for the remaining users. We relax these two sets of constraints to get sufficient conditions for the equilibrium rates and demonstrate two special cases when the rates obey these extra constraints for free, thus giving us real equilibria.

Assume for the moment that we know the set of users who have positive probability of being engaged. This set of users is the equivalent of the matching for the Followership Game. We will construct a candidate Nash Equilibria based on such a set. Assume that F'_i is this set of users who follow celebrity i who have positive probability of being engaged; i.e., $F'_i = \{j \in F_i \mid \sum_{l \in C_j} r_l \leq 1\}$. Given F'_i and relaxing the constraints, the optimal rate r_i^* is

$$r_i^* = \arg \max_{r_i} r_i \sum_{j \in F'_i} (1 - \sum_{l \in C_j} r_l).$$

Using the first order conditions we get

$$r_i^* = \frac{|F'_i| - \sum_{j \in F'_i} \sum_{l \in C_j - \{i\}} r_l}{2|F'_i|} \quad (1)$$

which is linear with respect to the rates. Therefore, if the set of users who have positive probability of being engaged is given we can solve the corresponding set of linear equations and calculate r_i^* 's for all celebrities. However, we then require that

$$\forall i, j : j \in F_i - F'_i \Rightarrow \sum_{l \in C_j} r_l^* \geq 1 \quad (2)$$

$$\forall i, j : j \in F'_i \Rightarrow \sum_{l \in C_j} r_l^* \leq 1. \quad (3)$$

Conditions 2 and 3 imply that the solution to the linear system is self-consistent with the assumption about F'_i , which users have positive probability of being engaged. Further, condition 3 demands that the rates are non-negative in combination with the linear system⁵.

Like in Section 2.3, a solution to the above linear system that obeys the conditions is not necessarily a Nash Equilibria. It is then necessary to check to ensure that no celebrity has incentive to unilaterally deviate their rate and change the set of possibly engaged users.

In section 3.1, we prove the existence of pure strategy Nash equilibrium for two classes of sparse graphs and regular graphs by computing a solution to the linear system and proving that no unilateral deviation is profitable.

3.1 Pure equilibria under sparse or regular users

Proposition 5. *A pure strategy Nash equilibrium exists for the engagement game when all users have the same degree.*

Proof. We show that if the graph is d -regular, i.e., $|C_j| = d$ for all j , every celebrity tweeting at rate $1/(d+1)$ is a pure strategy Nash equilibrium of the game. We let $F'_i = F_i$ and solve the linear system defined by equality 1 for every i . The solution is $r_i^* = 1/(d+1)$. We now confirm that the solution satisfies conditions 2 and 3. Condition 2 is obviously satisfied since $F_i - F'_i = \emptyset$; since $\sum_{l \in C_j} r_l^* = d/(d+1)$, condition 3 is also satisfied. Now let's examine whether a particular celebrity

⁵ As mentioned above, these conditions will not be necessarily satisfied for every choice of F'_i 's, and in fact, it is not clear even if there exists any choice of F'_i 's for which Equation 1 and conditions 2 and 3 all hold. We conjecture that if a Nash Equilibrium exists, then there must exist such a self-consistent choice of F'_i 's.

can profitably change their rate. All users are at $(d-1)/(d+1)$ without including this celebrity's rate. So this celebrity can either possibly engage their neighboring users with rate less than or equal to $2/(d+1)$ or cause them all to disengage with rate greater than $2/(d+1)$. They will always do the first, meaning that they will not deviate from the solution by causing a different engagement set. Hence, by equality 1, we obtain a pure strategy Nash equilibrium of the game.

Proposition 6. *A pure strategy Nash equilibrium exists for the engagement game when all users follow one or two celebrities.*

Proof. This is the case when $|C_j| \leq 2$ for all j . First note that since r_i^* is upper-bounded by $1/2$ for any set of possibly engaged users and each user follows at most two celebrities, every user is always possibly engaged and we do not need to consider unilateral deviations that change the set of possibly engaged users. Then letting $F'_i = F_i$ satisfies conditions 2 and 3 for any solution r_i^* . Therefore, any solution that satisfies equation 1 is a pure strategy Nash equilibrium of the game.

We now prove the existence and uniqueness of a solution to the linear system defined by 1 assuming that every celebrity has at least one follower. Equation 1 reduces to

$$r_i^* = \frac{|F'_i| - \sum_{l=1}^n a_{il}r_l^*}{2|F'_i|}$$

where a_{il} denotes the number of users that follow both celebrities i and l with $a_{ii} = 0$. Note that the set of followers of celebrity i consists of those who only follow celebrity i which we call M_i , and those who follow celebrity i and some other celebrity l . Define $n \times n$ matrix B such that $b_{ij} = a_{ij}$ for $i \neq j$, and $b_{ii} = 2|F'_i|$. Notice that the linear system defined by 1 has a unique solution if and only if the determinant of B is non-zero. Since $|F'_i| = |M_i| + \sum_{l=1}^n a_{il}$, every element on the diagonal of B is at least twice of the sum of all other elements in its column. Therefore, B is strictly diagonally dominant, and consequently, by Gershgorin circle theorem [8], the determinant of B is non-zero.

4 Conclusions

We have introduced new models of rate competition in asymmetric follow networks like Twitter. We hope our models will be further investigated for both empirical validation as well as further theoretical study. We also hope that our insights will highlight the importance of update rate to users of social media as well as that of effective filtering techniques to better automatically cope with such overload.

Several directions for future work present themselves.

One important open problem is to extend our followership model to non-bipartite networks where nodes are both users and celebrities and there is a correlation between the production and consumption utilities of agents. Another similar direction is to extend the engagement model to the non-bipartite case.

Another avenue for future work in this area relate to computing the price of anarchy and the price of stability even in the special cases when there are pure strategy equilibria for these models.

A natural alternate method of carrying out an analysis of information overload, particularly in the context of a fast growing network like Twitter is to postulate a stochastically growing model and study the results of a best response adjustment of rates during this stochastic growth. We could then try to find limiting properties of the network formed such as its degree sequence and the effect of the entry time of celebrities as well as their overall aggregate quality, and see if this provides a better model of network evolution than other more mechanistic and rate-oblivious models of network growth.

References

1. The Blog Herald, Twitter's Meteoric Rise, <http://www.blogherald.com/2010/06/28/twitters-meteoric-rise-compared-to-facebook-infographic/twitter-statistics-infographic-911/>
2. Eppler, Martin J. and Mengis, Jeanne. *The Concept of Information Overload: A Review of Literature from Organization Science, Accounting, Marketing, MIS, and Related Disciplines*, The Information Society, 20(5), 325-344 (2004)
3. Huberman, Bernardo, Romero, Daniel, and Wu, Fang. *Social networks that matter: Twitter under the microscope*. First Monday, 14 (1), 2009.
4. Matthew O. Jackson, *Social and Economic Networks*, Princeton University Press (2009).
5. Michael Sattinger, *Value of an Additional Firm in Monopolistic Competition*, The Review of Economic Studies, 51(2), 321-332 (1984).
6. Clay Shirkey on information overload versus filter failure, video from Web 2.0 Expo NY. <http://www.boingboing.net/2010/01/31/clay-shirky-on-infor.html> (2010).
7. P.G. Doyle and J.L. Snell. *Random walks and electric networks*. The Mathematical Association of America, 1984.
8. Horn, R.A. and Johnson, C.R. *Matrix analysis*. Cambridge University Press, 1990.