

The Role of Organization Hierarchy in Technology Adoption at the Workplace

Charalampos Chelmis

Department of Computer Science
University of Southern California
Los Angeles, USA
chelmis@usc.edu

Viktor K. Prasanna

Ming Hsieh Department of Electrical Engineering
University of Southern California
Los Angeles, USA
prasanna@usc.edu

Abstract—Popular social networking sites have revolutionized the way people interact on the Web, enabling rapid information dissemination and search. In an enterprise, understanding how information flows within and between organizational levels and business units is of great importance. Despite numerous studies in information diffusion in online social networks, little is known about factors that affect the dynamics of technological adoption at the workplace. Here, we address this problem, by examining the impact of organizational hierarchy in adopting new technologies in the enterprise. Our study suggests that middle-level managers are more successful in influencing employees into adopting a new microblogging service. Further, we reveal two distinct patterns of peer pressure, based on which employees are not only more likely to adopt the service, but the rate at which they do so quickens as the popularity of the new technology increases. We integrate our findings into two intuitive, realistic agent-based computational models that capture the dynamics of adoption at both microscopic and macroscopic levels. We evaluate our models in a real-world dataset we collected from a multinational Fortune 500 company. Prediction results show that our models provide great improvements over commonly used diffusion models. Our findings provide significant insights to managers seeking to realize the dynamics of adoption of new technologies in their company, and could assist in designing better strategies for rapid and efficient technology adoption and information dissemination at the workplace.

Keywords—agent based computational models; adoption dynamics; influence; evolutionary models; diffusion of innovations

I. INTRODUCTION

Researchers have well studied the importance of social networks on information spread [1], [2], emphasizing particularly on information dissemination. Traditionally, diffusion and cascading behavior have been formalized as transmission of infectious agents in a population, where each individual is either infected or susceptible, and infected nodes spread the contagion along the edges of the network. There are, however, differences between information flows and the spread of viruses. While viruses tend to be indiscriminate, infecting any susceptible individual, information is selective and passed by its host only to individuals the host thinks would be interested in it. Diffusion models heavily rely on the premise that contagion propagates over an implicit network, the structure of which is assumed to be sufficient to explain the observed behavior. However, the structure of the underlying network has to be learned [3] from a plethora of historical evidence, i.e. cascades. Although diffusion theory brings up the importance

of friendship relations, adoption behavior is instead examined on the premises of the behavior of the entire population [2].

Unlike in online social networking sites where users create links to others who are similar to them [4], or whose contributions they find interesting [5], in a corporate environment, employees form “bonds” not because of similar “tastes” but due to tasks at hand or because of reporting-to relationships, i.e. organizational hierarchy. In this sense, there is no explicit “social network”, however, formal structures such as the organizational hierarchy may provide hints of influence at the workplace. The dynamics of information diffusion on a corporate environment are yet unknown and may be entirely different from online social networks. The interplay between formal structure and information propagation at the enterprise has been recently examined [6]. The authors found that social and organizational structure significantly impacts the spreading process of emails, while at the same time indicating context independence. In our study, on the contrary, we do not know the chain of infections, i.e. we do not observe who influences whom. Instead, we empirically quantify the role of reporting-to relationships and local behavior (teammates), as well as the effect of global influence (overall popularity) in the spread of technology adoption at the workplace.

To characterize the adoption mechanism of new technologies at the workplace, we propose two simple and intuitive agent-based computational models with the least possible number of parameters. We emphasize on accurately modeling the cumulative number of adoptions over time, rather than trying to predict which node in the network will infect which other nodes. In this sense, we not only model the influence each node has on the diffusion (*microscopic modeling*), permitting user behavior to vary according to the behavior of the general crowd, but we also provide a simple mechanism by which adoption rate rises and decays over time (*macroscopic dynamics*). For our study, we have acquired the organizational hierarchy of a Fortune 500 multinational company. In addition, we gathered adoption logs of the internal microblogging service, which resembles Twitter, during the first two years of adoption of the service in the enterprise. This dataset allows us to empirically characterize individual dynamics and influence, and examine the spread of adoption through the hierarchy.

The rest of the paper is organized as follows. We describe our dataset in Section II. We study the impact of hierarchical structure on the way adoption spreads in Section III, and we examine employees behavior with respect to overall popularity

of the microblogging service in Section IV. To capture the macroscopic, temporal dynamics of adoption at the workplace, we propose two novel models that effectively model user behavior with respect to the entire population and individual influence in Section V. In Section VI, we provide extended social simulation results of our agent-based computational models of adoption at the workplace. We provide an overview of the most relevant related work which has been undertaken in this area in Section VII. Finally, we discuss the findings of our work and draw our conclusions in Section VIII.

II. DATASET

The company we studied is a Fortune 500 multinational company, which operates outside the IT-sector. Our dataset consists of a snapshot of the organizational hierarchy, containing over $12K$ employees. Our dataset further contains employees' join logs during the first two years of adoption of a microblogging service from the enterprise (July 2, 2010 to March 22, 2012). During this time period, the number of employees who join the service increases dramatically. Even though, not all employees have joined the microblogging service by the time we obtained the raw data for this paper, a broad spectrum of employees (9,421 users) had joined the microblogging service (77.35% of hierarchy dataset), sharing 19,371 status updates and exchanging 20,370 replies [7]. The functionality of the microblogging service resembles that of Twitter, imposing no restrictions on the way people interact or who they chose to follow. As in Twitter, users author messages in the enterprise microblogging service, and form threaded discussions. The main purpose of the corporate microblogging service is to promote and enable collaboration and sharing within the enterprise. The ultimate goal of the corporate microblogging service is to become the primary platform for asynchronous collaboration and colleagues' communication.

The company did not officially initiate usage of the microblogging service. Rather, it was independently initiated by an employee, in the begging of July, 2010. It was not promoted or even mentioned in any formal corporate communications. Our dataset does not contain information with respect to growth and invitations. We can only speculate that growth was achieved through email and word of mouth invitations. More details on the topological properties of the corporate microblogging service, its dynamics and characteristics, and the interplay between its social and topical components, users' homophily and activity, as well as latent topical similarity and link probability can be found at [7].

III. EFFECT OF ORGANIZATION HIERARCHY

The underlying process of influencing employees towards adopting the microblogging service is unknown and non trivial. Here, we assume that when an employee chooses to join the corporate microblogging service, she then has some influence on the employees who directly report to her, according to the formal organizational chart. Some of these employees will choose to join, which will in turn influence some of their team members into joining themselves and so on. Therefore, we assume that an employee's decision to join depends on: 1) direct influence by her manager, 2) peer influence by her teammates, and 3) social influence resulting from the overall popularity of the microblogging service in the enterprise. In

this section we seek supporting evidence on the influence inflicted by managers to employees reporting directly to them.

Assume that manager u urges her team members to join the microblogging service. A directed link e_{ju} exists if employee j directly reports to u according to the formal organizational hierarchy. If j joins the microblogging service after u , we call her join an "influenced join". We counted the number of employees who joined the microblogging service after their manager and found that there are three classes of employees: (i) employees who did *not* join the microblogging service even if their manager did (10.94%), (ii) employees who *did join* the microblogging service *before* their manager (36.04%), and (iii) employees who *did join* the microblogging service *after* their manager (53.01%).

Let N be the total number of employees directly reporting to manager u . Let K be the number of employees in N that joined the microblogging service after their manager u , and k be the total number of employees in N that joined the microblogging service after their manager u within the first n draws. The stochastic process according to which employees directly reporting to u choose to join the microblogging service is described by the "urn model" [8], in which n balls are drawn without replacement from an urn containing N balls in total, of which K are white. The probability $P(X = k|K, N, n)$ that k of the first n employees reporting to manager u , joined the microblogging service after their manager purely by chance is equivalent to the probability that k of the n balls drawn from the urn are white. We set $n = 8$, calculating the number of employees that joined the microblogging service after their manager within the first 8 draws. This probability is given by the hypergeometric distribution:

$$P(X = k|K, N, n) = \frac{\binom{K}{k} \binom{N-K}{n-k}}{\binom{N}{n}}. \quad (1)$$

We plot the average number of employees that joined the microblogging service after their manager during the first n samples as a function of the number of employees that joined the microblogging service after their manager. Figure 1 shows the result. The scatter plot is approximated [8] by the Weibull cumulative distribution ($\hat{k} = 24(1 - e^{-(0.02K)^{0.84}}$). We use this expression to estimate the expected number \hat{k} of employees to join the microblogging service after their manager within the first n joins for a manager with K employees reporting to her that joined the microblogging service after her. Using Equation 1, we calculate the probability that \hat{k} employees joined after their manager purely by chance. We found that for $K > 3$, this probability is exceedingly small. Since it is exceedingly highly unlikely for employees to adopt the microblogging service after their manager purely by chance, we conclude that the number of employees who joined after their manager u is a prominent indicator of u 's influence.

A. Influence Score

Let N_j denote the number of employees who directly report to u and have joined the microblogging service. Let $\alpha \leq N_j$ be the number of employees that report to u and have joined the microblogging service after u , and let $q \leq N_j$ be the number of employees that report to u and have joined the microblogging service before u . While a high number of employees reporting




Fig. 1: Average number k of employees that joined the microblogging service after their manager, within the first n samples vs the total number K of employees that joined the microblogging service after their manager, and approximation.

to u that have joined the microblogging service after u implies that u has high influence, a high q value is an indicator that one lacks influence. We propose an adaptation of the z-score [9], as a measure that combines the number of employees that have joined before and after their supervisor. Influence score (“ ι -score”) measures how different this behavior is from a user with “random” influence, i.e. a manager the employees reporting to whom join after him with probability $p = 0.5$ and before him with probability $(1 - p) = 0.5$. We would expect such a random influencer to have $N_j * p = N_j/2$ team-members who joined after their supervisor with a standard deviation of $\sqrt{N_j * p * (1 - p)} = \sqrt{N_j}/2$ [9]. The ι -score measures how many standard deviations above or below the expected “random” value a manager u lies:

$$\iota(u) = \frac{\alpha - N_j/2}{\sqrt{N_j}/2} = \frac{\alpha - q}{\sqrt{\alpha + q}}. \quad (2)$$

If the employees reporting to manager u have joined the microblogging service after u about half of the time, u 's ι -score will be close to 0. If they join after u more often than not, u 's ι -score will be positive, otherwise, negative. We also calculate the time-independent ι -score of employees using Equation 2, with the difference that $\alpha \leq N$ is the number of employees that have joined the microblogging service (time invariantly) and $q \leq N$ is the number of employees that have not joined the microblogging service. Above, we measured influence at the level of individual employees, assuming that influence scores are fixed in time, but that they differ from employee to employee. A more sophisticated model of influence might include some small increase (similarly for decrease) in influence score as a function of time. We stick to the simpler model for simplicity, and because our fundamental result is not sensitive to such details.

Next, we examine the correlation between ι -score of managers and the number of employees reporting to them (team size), hoping to get a clearer picture of the relationship between the two quantities. We characterize the average ι -score of managers with λ employees reporting to them as $\iota(\lambda) = \frac{1}{|u:\lambda_u=\lambda|} = \sum_{u:\lambda_u=\lambda} \iota(u)$. Figure 2a shows the average ι -score of managers with λ employees reporting to them, that

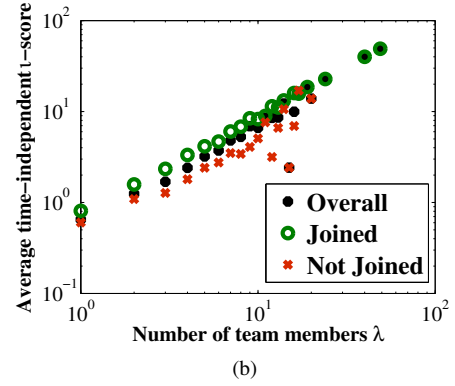
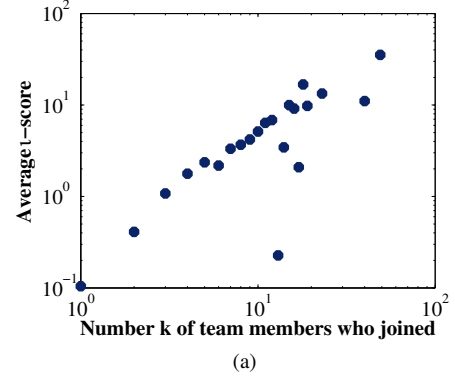


Fig. 2: (a) Average ι -score of managers with λ team members that have joined the microblogging service. (b) Average time-invariant influence of managers, who have themselves joined the microblogging service (similarly for those who have not joined), with λ team members.

have joined the microblogging service. Here, we focus on managers that have themselves joined the microblogging service, so that a time comparison of joining times is meaningful. A clear increasing trend is evident, providing a supporting evidence on top-down influential flow through the formal organizational hierarchy. Figure 2b shows the average time-independent ι -score of managers with λ employees reporting to them. Figure 2b further shows different plots of the average time-independent ι -score of managers based on the premise that they have joined the microblogging service themselves or not. The average time-independent ι -score of managers that have not joined the microblogging service exhibits more fluctuations due to greater data sparsity. In every case, the average time-independent ι -score of managers that have joined the microblogging service is slightly higher than for managers that have not joined the service. Even though we cannot at the time explain the reasons why this effect appears, the average time-independent ι -score increases for both classes as the team size λ increases, clearly indicating a strong correlation between the two quantities. We explain this trend as a prominent indicator of influence imposed by managers to employees reporting directly to them.

We now turn our attention to the impact of organizational levels. Here, we assume that influence scores are characteristic of a particular level at the organization hierarchy tree, are fixed in time, and are the same for all employees at that particular

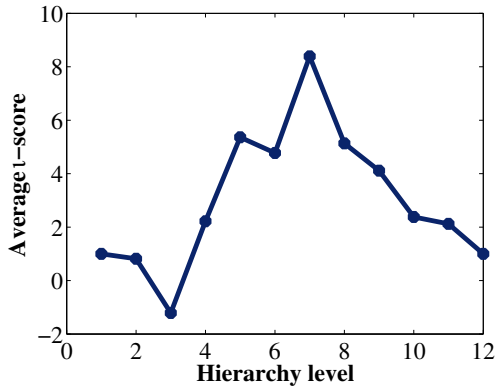


Fig. 3: Average influence score as a function of hierarchy level.

level. To compute the average influence score for hierarchy level l , we first find employees m that belong to level l . We then compute the total number of employees N that directly report to managers in E_l . Quantities α and q are defined as before, with the difference that they now operate on the total number of employees N that directly report to managers in E_l . Finally, we use Equation 2 to calculate the influence score for each level. Levels are ascending from the CEO (level 1) to lower levels. Level 13, which represents bottom level employees in our dataset, contains employees with no team members reporting to them.

Figure 3 shows the results. Level 13 has no influence score, thus it does not appear in Figure 3. Most levels exhibit positive influence scores, with the exception of higher levels, that are closest to the CEO. Particularly, level 3, exhibits negative influence on average. As before, we measured influence at the granularity of hierarchical levels, assuming that influence scores are fixed in time, but that they differ from level to level. A more sophisticated model of influence might include some small increase (similarly for decrease) in influence score as a function of time, and also introduce a balancing factor based on the number of total employees at a level and the number of total employees reporting to them. While it is intuitive to assume that higher levels in the organization would have higher impact due to the report-to relationships involved, our study suggests that middle levels are more successful in influencing employees lying lower in the hierarchy. Even though we do not have supporting evidence from other use-cases, we conjecture that middle-level managers are the most influential with respect to “convincing” others to adopt new technologies (in this case the new microblogging service).

IV. EFFECT OF PEER PRESSURE

We study the problem of progressive diffusion, where the employees who adopt the microblogging service become infected and do not become healthy again (i.e. employees do not unsubscribe the service once they join). Classic models of social and biological contagion (e.g. [10], [11]) and observational studies of online contagion [12], [13], [14], [15] predict that the likelihood of infection increases with the number of infected contacts. However, recent studies suggest that this correlation can have multiple causes that might be unrelated to social influence processes [16]. In our observational study of

microblogging service adoption at the workplace, this assumption suggests two alternative modeling scenarios. According to the first scenario, an employee is more likely to adopt the microblogging service if more of her teammates join the service (Section IV-A). According to the second scenario, an employee is more likely to adopt the microblogging service as its popularity increases (Section IV-B). Our goal in this section is to find models that will provide a good fit with respect to the probability of adoption for each user given the actions of their teammates (local neighborhood) or overall popularity (global influence).

A. Independent Cascade Model

Influence of friends is generally modeled to be additive. For instance, the independent cascade model (ICM) [17] states that a node has n independent chances to become infected, where n is the number of infected “friends”. In our case, every node can be infected only once, and once infected, it stays infected. Because of the structure of the organizational hierarchy, employee u ’s “friends” may include either (i) her teammates alone, or (ii) her teammates and her direct supervisor. Starting with a single employee who has joined the microblogging service, employees *susceptible* to infection, decide to join the microblogging service with some probability that depends on the number of their infected “friends”. We model the influence employees receive by their “friends” as multiple exposures to an infection according to ICM [17] as $p_{ICM} = 1 - (1 - \lambda)^n$.

We measured this quantity on our dataset, by isolating the employees in two classes: a) those who had exactly n “friends” joining the microblogging service and did not join, and b) those who had exactly n “friends” joining the microblogging service before they themselves joined. We found that the likelihood of adoption when no “friends” have joined is remarkably high (0.7581 when considering teammates only and 0.6807 when the supervisor is also considered). In both cases, the likelihood of adoption becomes 1 when at least one “friend” has joined the service. We conclude that the relationship between the number of “friends” that have joined and likelihood of joining most probably reflects heterogeneous popularity of the microblogging service across teams [16]. Therefore, the naive conditional probability does not directly give the probability increase due to influence via multiple joining “friends” [16].

B. Exponential Growth Model

We studied earlier the effect of multiple teammates and neighbors of an employee u on the probability of u to join the microblogging service. Even though we discovered a positive correlation, we argued that this correlation might be an effect of multiple causes. We hypothesized that the more popular the microblogging service is for a team, the more likely it is for multiple team members to adopt it. Further, as employees observe others adopting the microblogging service, they may not only be more likely to adopt the service, but the rate at which they do so may quicken as the popularity of the service increases. In this section, we venture to explore this hypothesis.

Figure 4 shows the probability that an employee will join the microblogging service as a function of the service popularity. Intuitively, as more people adopt the microblogging service, a certain “buzz” around the service begins to unfold,

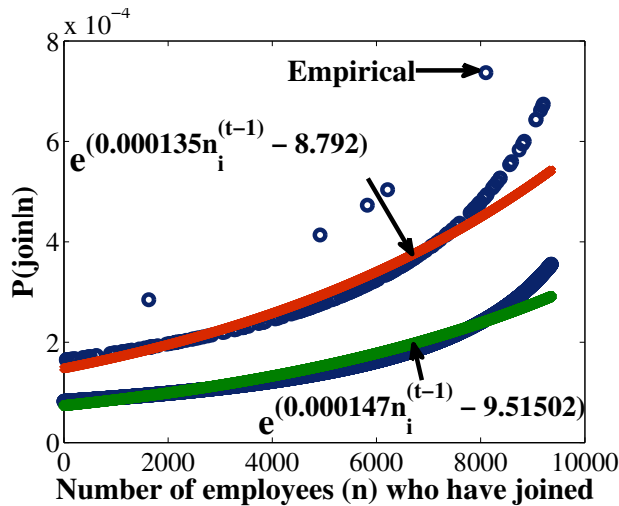


Fig. 4: Probability an employee joins the microblogging service given that n employees have adopted the service before. Solid lines depict probability estimates calculated with the exponential growth model.

increasing the probability of others joining the service as well. Interestingly, Figure 4 reveals that employees join the microblogging service following two very different, clearly distinctive patterns. According to the “*optimistic*” pattern (red line), the probability of adoption increases more profoundly as the overall number of people who join increases. Contrary, the “*pessimistic*” pattern (green line) yields a probability of adoption that increases marginally as the total number of people who join the service increases over time. Even though we cannot at the time explain this effect, these two distinct classes of people remain to be fully understood in future work, in conjunction to surveys and targeted interviews.

V. MODELING ADOPTION

What is the underlying hidden process that drives adoption of new technologies at the workplace? Our goal here is to find a generative model that generates the observed adoption process of the new microblogging service at the enterprise we are studying, given the organizational hierarchy. Such a model should exhibit the properties we observed in Sections III and IV and reproduce the true cumulative number of adoptions. We aim for simple and intuitive modeling with the least possible number of parameters.

Prior work on modeling complex networks in social, biological and technological domains has focused on replicating one or more aggregate characteristics of real world networks [18]. Here, we take a different approach. Instead of having a target network to generate, we let individual influence and peer pressure dynamics determine the diffusion process of adoption of the new microblogging service over the formal organization hierarchy. We propose two models that account for influence effects imposed by the formal organizational structure. We compare our results to the true epidemic and we show that the estimates produced by our models are consistent with the real observations.

A. Complex Contagion Model

From the empirical analysis presented in the previous sections, we incorporate the following dynamics into our model:

- Employees are influenced by their managers to join the microblogging service.
- Employees have multiple chances to get infected (join). Once an employee is infected, she cannot recover, i.e. an employee does not unsubscribe from the service.
- As employees observe others adopting the microblogging service, they are not only more likely to adopt the service, but the rate at which they do so quickens as the popularity of the service increases.

We begin by selecting a single node from the organization hierarchy to start the infection. We chose the seed node to be the exact employee that first registered to the microblogging service according to our dataset. At each time step, the virus can be spread as follows. Each node that was infected at time $t - 1$ has n chances to infect the n employees that directly report to her, each with probability p , at time t . Once a node is infected, it cannot be infected again. An infected employee is not allowed to infect her direct supervisor, so following this strategy, the virus can only propagate towards the leaves of the hierarchy tree. Once all infected nodes are examined, healthy nodes have the chance to be “randomly” infected by observing the general popularity of the microblogging service up to time $t - 1$. For n_i^{t-1} total infected nodes at time $t - 1$, the probability of “random” infection is computed using the pessimistic exponential growth pattern fit ($p_{EG} = e^{(0.000147n_i^{t-1} - 9.51502)}$) from Section IV-B. Note that the selection of the pessimistic exponential growth pattern is a conservative choice in that it does not unfairly help our model in predicting the cumulative number of adoptions over time.

B. Complex Cascade Model

The model we described above spreads the adoption of the microblogging service over the formal organization hierarchy as a virus, which leaves a trail whenever employees are infected by their supervisors. To model this we used parameter p , which measures how infectious supervisors are, and parameter p_{EG} that controls the effect of overall growing popularity of the microblogging service over time. Here we take an alternate approach based on which, nodes choose to become infected after examining their immediate neighborhood (which includes both the supervisor and employees directly reporting to them) or after examining the overall growing popularity of the microblogging service over time.

We start with the organization hierarchy, and two colors. Let red represent employees who have joined the microblogging service and blue those that have not. We choose a single node to be the seed user, i.e. have color red. All other users are painted blue. As before, we chose the seed node to be the exact employee that first registered to the microblogging service according to our dataset. At each time step, nodes painted blue (not infected), calculate the payoff of picking the color red over blue, and decide their color $f(\text{color})$ as follows:

$$f(\text{color}) = \begin{cases} \text{red}, & \alpha \frac{n_{\text{red}}}{n} > \beta \frac{n_{\text{blue}}}{n} \\ \text{blue}, & \text{otherwise} \end{cases}, \quad (3)$$

where n_{blue} denotes the number of blue neighbors, n_{red} denotes the number of red neighbors and $n = n_{blue} + n_{red}$ is total number of neighbors. Parameters α and $\beta = 1 - \alpha$ denote the rewards for choosing red and blue accordingly. Once a node is painted red, it cannot change color again. Finally, nodes have the chance to be “randomly” infected by observing the general popularity of the microblogging service up to time $t - 1$. As in our contagion model, for n_i^{t-1} infected nodes, the probability of “random” infection is computed using the pessimistic exponential growth model fit ($p_{EG} = e^{(0.000147n_i^{t-1} - 9.51502)}$) from Section IV-B.

VI. EVALUATION

In this section, we validate our models by extensive numerical simulations. We begin with the organization hierarchy of 12,170 employees, and infect the true initiator of the epidemic (the employee who first joined the microblogging service). Each time step represents a day. We let our models run for 600 steps, or until all employees are infected. We compare the obtained epidemics against the real cumulative number of adoptions extracted from our dataset. We experimented with various values of infection probability for our contagion model and parameters α and β for our cascade model. In the end, we decided to use $p = 0.3$ for our contagion model, and $\alpha = 0.82$ and $\beta = 0.18$ in our cascade model. We simulated our models 10 times and report our findings. We compare three properties of the simulated epidemics as opposed to the true number of adoptions over time: (i) overall number of infections, (ii) cumulative number of infections over time, and (iii) total time required to infect N employees. We find that our models’ estimates are consistent with the real observations.

A. Baselines

We compare our proposed models’ ability to approximate the true cumulative distribution of infected users with three models, which have shown superior performance in the task of information and innovation diffusion in social networks.

- **Susceptible-Infected Model (SI)** [19]: According to the SI model, each node can infect her neighbors, each with probability p_{SI} . We considered the Susceptible-Infected-Susceptible (SIS) and Susceptible-Infected-Resistant (SIR) models [20], as well as the Susceptible-Infected-Dead (SID) model [21] as alternatives to model social contagion, as these models are widely used in prior work. These models however do not appropriately capture the semantics of adoption, according to which, an employee that joins the microblogging service does not unsubscribe, thus returning to the susceptible state, or becoming resistant. Further, our analysis did not provide any supporting evidence for the hypothesis that infected employees do not infect others, thus modeling them as “dead” is not appropriate in this case.
- **Independent Cascade Model** [17] (see Section IV).
- **Diffusion Model (DM)** [22], [23], [24]: Each individual’s willingness to adopt the microblogging service at time t , U_u^t , is modeled by three main elements: the service’s stand-alone benefit, network effects, and

the idiosyncratic reservation utility. Formally, $U_u^t = Q_u + \gamma N_u^{(t-1)} - R_u$, where, Q_u represents the service’s intrinsic value perceived by employee u , which is not affected by whether other people adopt it or not. $N_u^{(t-1)}$ represents the proportion of adopters in u ’s neighborhood at time $t - 1$, and γ denotes the relative importance against stand-alone benefits. R_u indicates u ’s inherent reluctance or reservation about adopting the new service.

B. Experimental Results

First, we study simulation results produced by the baselines, i.e. the SI, ICM and DM models. For brevity, we focus on reporting results for the SI model only. The analysis of ICM and DM simulations yielded analogous results, therefore we believe our conclusions to be robust. Figure 5a shows the true user adoption curve, compared to simulation results produced by the SI model, for varying infection probability values. We notice that simulation models do not fit the real cumulative number of adoptions over time. High infection probability values result in sudden outbreaks, whereas very small probability values result in smooth cumulative distributions that do not exhibit the statistical properties of the true cumulative number of infected users. The total number of infections and the time required to infect the whole body of employees is also inconsistent with the observed adoption curve.

Next, we show the outcome of ten runs of our complex contagion model (see Section V-A) in Figure 5b. The figure also shows the average of the ten runs. Notice a very good alignment between the reality and simulated epidemics in all cases. Not all runs result in the total number of true infections by the time threshold. Further, a few runs overestimate the cumulative number of infections, resulting in rapid epidemics. Unlike the baselines, our complex contagion model fits more naturally the true cumulative number of infected users in all cases. Particularly, the simulation results remarkably follow the speedups and slowdowns of adoption over time, exhibiting non-linear characteristics as the true adoption curve. Some runs diverge from the true curve after about 400 days. However, running the model numerous times and averaging the results seems to adequately approximate the statistical properties of the true cumulative number of infected users. We conclude that this is a direct result of the asymmetric contagion due to the hierarchical influence to adoption and the integration of peer pressure due to growing popularity of the microblogging service at the enterprise.

Finally, we present the outcome of ten runs of our complex cascade model (see Section V-B), and their average, in Figure 5c. In this case too, simulated epidemics match the reality very well. Similarly to epidemics produced by our cascade model, not all runs result in the total number of true infections by the time threshold. Further, smooth regimes of adoption, speedups and slowdowns of the acceptance of the microblogging service from employees is apparent. Unlike our cascade model, this model slightly overestimates the cumulative number of infections. In all cases however, we find that this model too fits rather closely to the true cumulative number of infected users, replicating the statistical properties of the empirical epidemic.




Fig. 5: True and predicted cumulative number of employees who have adopted the microblogging service (i.e. infected users). Time is measured in days. Solid line curves represent the outcome of (a) the SI model for various probabilities of infection, (b) ten runs of our complex contagion model (see Section V-A), and (c) ten runs of our complex cascade model (see Section V-B).

VII. RELATED WORK

The importance of social networks in information dissemination has been thoroughly investigated [3], [16], [2]. In online social networks in particular, where individuals tend to organize into groups based on their common activities and interests (a phenomenon known as homophily [4]), it has been hypothesized that the network structure (friendship or interaction) affects the way information spreads, and that adoption quickens as the number of adopting friends increases [13]. However, many times a node activation is not just a function of the social network but also depends on many other factors like imitation [1]. This has led to the development of epidemiology models [20] and computational approaches that are based on thresholds models [10], deterministic or stochastic [25]. Each agent has a threshold that, when exceeded, leads the agent to adopt an activity. When the threshold is applied within a local neighborhood [26], [27], local models emerge [17]. Instead, global diffusion models perform thresholding to the whole population [2].

Diffusion models heavily rely on the premise that contagion propagates over an implicit network, which has to be learned from a plethora of historical evidence, i.e. cascades. User characteristics such as topical or latent interests have to be considered in user-to-user content transfers, whereas users' homophily shapes the structure of the network through which information flows. In a corporate environment, employees form "bonds" not because of similar "tastes" but due to a task at hand (i.e. a function to be completed or an organizational need) or because of reporting-to relationships (i.e. team members reporting to their supervisor). [3] examined the problem of inferring the unobserved directed network over which cascades propagate in online social networks. Unlike their approach, which requires traces of numerous different explicit cascades to be given as inputs, we solely rely on one *implicit* sample to infer influence between employees at the workplace. In fact, many influence models have been proposed to rank actors within a social network [8]. However, the underlying dynamic process occurring on the network may not be applicable to the organizational hierarchy. Influence models typically do not take the topology of the network into account, and when they do, they make assumptions about the details of the underlying dynamic process taking place on the network.

In our empirical study, we characterize individual dynamics and influence, and examine the spread of adoption through the formal organizational hierarchy.

Even though most prior work has mainly focused on publicly available online social networks, microblogging capabilities have penetrated the enterprise as well [28]. Contrary to online social networks, microblogging services for enterprises are primarily designed to improve intra-firm transparency and knowledge sharing. However, the adoption of such collaborative environments presents certain challenges to enterprises [29]. [28] provided a case study on the perceived benefits of corporate microblogging and barriers to adoption. Key factors influencing microblogging systems adoption in the workplace include: privacy concerns, communication benefits, perceptions regarding signal-to-noise ratio, and codification effort, reputation, expected relationships, and collaborative norms [29]. The work, closest to ours, [6] examined email threads and the formal network (e.g. hierarchical structure) imposed by a large technology firm. They argued that the spreading process (to whom and how fast people forward information) can be well captured by a simple stochastic branching model. In our study, on the contrary, we do not know the chain of infections (i.e. we do not observe who influences whom). Instead, we use the outcome of our empirical study to quantify influence as a result of individual pressure from supervisors towards their team members, as well as an effect of global popularity.

VIII. CONCLUSION

In this paper, we studied the effect of the formal organizational structure, to the adoption mechanism of a microblogging service at the enterprise. We addressed the factors that govern the process of adoption at both microscopic and macroscopic levels. We found, microscopically, that employees' tendency towards adopting or not the new microblogging service is influenced by their direct supervisors (dependency on the network structure). We proposed ι -score as a prominent indicator of influence imposed by managers to their teams and we offered proof that middle level managers are on average more successful in promoting the adoption of the new service. Further, we empirically measured employees' likelihood of adopting the new microblogging service with respect to the behavior of the general crowd. We revealed two distinct

patterns, that capture the adoption likelihood increment as a function of the overall service popularity among the employee population. We incorporated our findings into two intuitive and simple adoption mechanisms, which capture both the local and global influence, accurately reproducing the adoption process at the macroscopic level. Prediction results show that our models provide great improvements over commonly used diffusion models. Our findings have important implications to enterprises' understanding of the mechanisms driving adoption of new technologies, and could assist in designing better strategies for rapid and efficient technology adoption and information dissemination within the corporation.

A limitation of our study is that we estimate causal effects only within the formal organizational chart, due to the fact that we are unable to observe the actual adoption "cascade" (i.e. who really influences whom). We are planning to further evaluate our results with extended surveys and targeted interviews, as well as incorporate more datasets in future work. We also plan to extend our models to allow for influence scores to vary over time, as well as incorporate different roles individual assume in the adoption process, accounting for influence variations as a function of employees' level in the organization hierarchy. We would also like to investigate the effect of network evolution (e.g. layoffs, or new hires) on influence, since one's influence may intuitively increase with seniority in the company. Finally, it would be interesting to study adoption dynamics in the presence of competing technologies.

ACKNOWLEDGMENT

We are thankful to Greg Harris and Daphney-Stavroura Zois for constructive feedback on this work. This work is supported by Chevron Corp. under the joint project, Center for Interactive Smart Oilfield Technologies (CiSoft), at the University of Southern California.

REFERENCES

- [1] J. Yang and J. Leskovec, "Modeling information diffusion in implicit networks," in *Proceedings of the 2010 IEEE International Conference on Data Mining*. Washington, DC, USA: IEEE Computer Society, 2010, pp. 599–608.
- [2] C. Budak, D. Agrawal, and A. El Abbadi, "Diffusion of information in social networks: Is it all local?" in *2012 IEEE 12th International Conference on Data Mining (ICDM)*, 2012, pp. 121–130.
- [3] M. Gomez Rodriguez, J. Leskovec, and A. Krause, "Inferring networks of diffusion and influence," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. New York, NY, USA: ACM, 2010, pp. 1019–1028.
- [4] M. McPherson, L. S. Lovin, and J. M. Cook, "Birds of a feather: Homophily in social networks," *Annual Review of Sociology*, vol. 27, no. 1, pp. 415–444, 2001.
- [5] J. Weng, E.-P. Lim, J. Jiang, and Q. He, "Twitterrank: finding topic-sensitive influential twitterers," in *Proceedings of the third ACM international conference on Web search and data mining*. New York, NY, USA: ACM, 2010, pp. 261–270.
- [6] D. Wang, Z. Wen, H. Tong, C.-Y. Lin, C. Song, and A.-L. Barabási, "Information spreading in context," in *Proceedings of the 20th international conference on World wide web*. New York, NY, USA: ACM, 2011, pp. 735–744.
- [7] C. Chelmiss and V. K. Prasanna, "Microblogging in the enterprise: A few comments are in order," in *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2012, pp. 62–70.
- [8] R. Ghosh and K. Lerman, "Predicting influential users in online social networks," in *Proceedings of KDD Workshop on Social Network Analysis*, 2010.
- [9] J. Zhang, M. S. Ackerman, and L. Adamic, "Expertise networks in online communities: structure and algorithms," in *Proceedings of the 16th international conference on World Wide Web*. New York, NY, USA: ACM, 2007, pp. 221–230.
- [10] M. Granovetter, "Threshold models of collective behavior," *American Journal of Sociology*, vol. 83, no. 6, pp. 1420–1443, 1978.
- [11] M. E. J. Newman, "Spread of epidemic disease on networks," *Physical Review E*, vol. 66, no. 1, p. 016128, Jul 2002.
- [12] A. Anagnostopoulos, R. Kumar, and M. Mahdian, "Influence and correlation in social networks," in *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. New York, NY, USA: ACM, 2008, pp. 7–15.
- [13] E. Bakshy, B. Karrer, and L. A. Adamic, "Social influence and the diffusion of user-created content," in *Proceedings of the 10th ACM conference on Electronic commerce*. New York, NY, USA: ACM, 2009, pp. 325–334.
- [14] M. Cha, A. Mislove, and K. P. Gummadi, "A measurement-driven analysis of information propagation in the flickr social network," in *Proceedings of the 18th international conference on World wide web*. New York, NY, USA: ACM, 2009, pp. 721–730.
- [15] J. Leskovec, L. A. Adamic, and B. A. Huberman, "The dynamics of viral marketing," in *Proceedings of the 7th ACM conference on Electronic commerce*. New York, NY, USA: ACM, 2006, pp. 228–237.
- [16] E. Bakshy, I. Rosenn, C. Marlow, and L. Adamic, "The role of social networks in information diffusion," in *Proceedings of the 21st international conference on World Wide Web*. New York, NY, USA: ACM, 2012, pp. 519–528.
- [17] D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the spread of influence through a social network," in *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. New York, NY, USA: ACM, 2003, pp. 137–146.
- [18] M. Newman, "The structure and function of complex networks," *SIAM Review*, vol. 45, no. 2, pp. 167–256, 2003.
- [19] J. A. Jacquez and C. P. Simon, "The stochastic si model with recruitment and deaths i. comparison with the closed sis model," *Mathematical Biosciences*, vol. 117, no. 1-2, pp. 77–125, 1993.
- [20] H. W. Hethcote, "The mathematics of infectious diseases," *SIAM Review*, vol. 42, no. 4, pp. 599–653, 2000.
- [21] C. Kamp, "Untangling the interplay between epidemic spread and transmission network dynamics," *PLoS Computational Biology*, vol. 6, no. 11, p. e1000984, 11 2010.
- [22] E. Abrahamson and L. Rosenkopf, "Social network effects on the extent of innovation diffusion: A computer simulation," *Organization Science*, vol. 8, no. 3, pp. 289–309, 1997.
- [23] F. M. Bass, "A new product growth for model consumer durables," *Manage. Sci.*, vol. 50, no. 12 Supplement, pp. 1825–1832, December 2004.
- [24] H. Choi, S.-H. Kim, and J. Lee, "Role of network structure and network effects in diffusion of innovations," *Industrial Marketing Management*, vol. 39, no. 1, pp. 170 – 177, 2010.
- [25] D. Strang and M. W. Macy, "In search of excellence: Fads, success stories, and adaptive emulation," *American Journal of Sociology*, vol. 107, pp. 147–182, 2001.
- [26] T. W. Valente, "Social network thresholds in the diffusion of innovations," *Social Networks*, vol. 18, no. 1, pp. 69–89, 1996.
- [27] S. Solomon, G. Weisbuch, L. d. Arcangelis, N. Jan, and D. Stauffer, "Social percolation models," *Physica A: Statistical Mechanics and its Applications*, vol. 277, no. 12, pp. 239 – 247, 2000.
- [28] J. Zhang, Y. Qu, J. Cody, and Y. Wu, "A case study of micro-blogging in the enterprise: use, value, and related issues," in *Proceedings of the 28th international conference on Human factors in computing systems*. New York, NY, USA: ACM, 2010, pp. 123–132.
- [29] O. Günther, H. Krasnova, D. Riehle, and V. Schöndienst, "Modeling micro-blogging adoption in the enterprise," *Proceedings of the 15th Americas Conference on Information Systems*, 2009.