

# Impact of Social Influence on Adoption Behavior: An Online Controlled Experimental Evaluation

Soumajyoti Sarkar, Ashkan Aleali, Paulo Shakarian  
Arizona State University, Tempe, USA  
Email: {ssarka18, aleali1, shak}@asu.edu

Mika Armenta, Danielle Sanchez, Kiran Lakkaraju  
Sandia National Laboratory, New Mexico, USA  
Email: {mlarmen, dnsanc, klakkar}@sandia.gov

**Abstract**—It is widely believed that the adoption behavior of a decision-maker in a social network is related to the number of signals it receives from its peers in the social network. It is unclear if these same principles hold when the “pattern” by which they receive these signals vary and when potential decisions have different utilities. To investigate that, we manipulate social signal exposure in an online controlled experiment with human participants. Specifically, we change the number of signals and the pattern through which participants receive them over time. We analyze its effect through a controlled game where each participant makes a decision to select one option when presented with six choices with differing utilities, with one choice having the most utility. We avoided network effects by holding the neighborhood network of the users constant. Over multiple rounds of the game, we observe the following: (1) even in the presence of monetary risks and previously acquired knowledge of the six choices, decision-makers tend to deviate from the obvious optimal decision when their peers make similar choices, (2) when the quantity of social signals vary over time, the probability that a participant selects the decision similar to the one reflected by the social signals and therefore being responsive to social influence does not necessarily correlate proportionally to the absolute quantity of signals and (3) an early subjugation to higher quantity of peer social signals turned out to be a more effective strategy of social influence when aggregated over the rounds.

## I. INTRODUCTION

The connections and interactions of individuals that comprise social networks are generally believed to impact decision-making in many domains including product selection [1], [4]. While there is general theoretical consensus that social influence, the phenomenon by which an individual’s opinions, behaviors, and decisions are influenced by other people [2], facilitates product selection, the empirical literature is actually quite torn. According to individual utility models, people adopt technologies when its benefits exceed its costs [3]. Because comparing every option can be cognitively costly and time consuming, individuals employ cognitive strategies and shortcuts that reduce the number of alternatives until one superior option is left [5]. Social signals factor into

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASONAM '19, August 27-30, 2019, Vancouver, Canada

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6868-1/19/08/\$15.00

<http://dx.doi.org/10.1145/3341161.3342882>

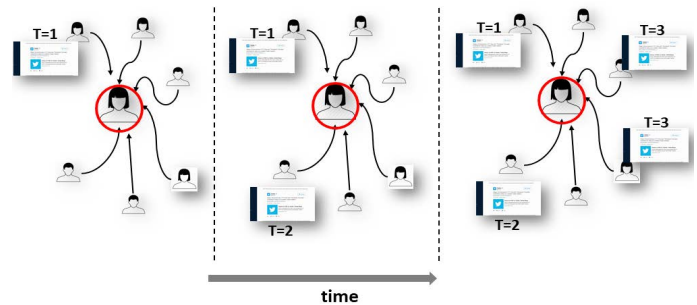


Fig. 1: This illustration demonstrates a *pattern of influence*. At the first timestep, one among 6 neighbors (followers) of a user shares a message, in the following timestep, another user shares the same message and in the final step, an 2 additional users. The user is thus exposed to a “pattern” of social signals we represent as  $V = \{1, 2, 4\}$  for that message.

the strategies because they provide a cost-efficient means of acquiring information [6]. Although social information is an additional layer that individuals must parse and weigh when making a decision, it is also a useful indicator of a product’s utility.

The literature documents several experimental results on the adoption of behaviors including network structure – such as the study conducted in [8]. Under this study, it was observed that individual adoption is much more likely when participants received social reinforcement from multiple neighbors in the social network as opposed to a single exposure. However as a major contribution, the research tests the effect of network structure on the dynamics of behavioral diffusion. Contrary to this, we quantify influence using only the number of signals temporally sent to a user irrespective of how the signals diffused to its neighbors prior to its own adoption. One of the main results from [8] show that the influence monotonically increases, although the increased likelihood of influence from  $k$  signals compared to  $k - 1$  signals peaks when  $k=2$ . Our results do not replicate this finding, as we did not consider network effects and the number of signals are time dependent and additionally, when each decision is supplemented by a reward.

In an observational study on the impact of repeated exposures on information spreading [9], the authors show that an overwhelming majority of message samples are more probable to be forwarded under repeated exposures, compared to those

under only a single exposure. Keeping these studies in mind, we develop an experimental framework to characterize the exposure effect under multiple signals but when the pattern of exposure could be controlled (we show an example of what a “pattern” is in Figure 1) - the experimental framework allows us to measure social influence through controlled exposures to neighborhood signals while avoiding confounding effects. It allows us to analyze how the signal versus timestep constraint impacts the selection choices of users. We use Amazon Mechanical Turk (AMT) to run an online, controlled decision-making game in which participants took on the role of a security officer at a bank. The subjects participated in several rounds of the game, where at each round, they selected one technology among 6 choices with differing utilities (only one among them was the optimal technology) - the subjects were able to view the choices made by their neighbors, with this modulation being part of the controlled setup that induces the exposure effect. Our main focus was to understand the effects of influence patterns on the choices made by individuals (whether individuals can be influenced to make a selection which inherently is not the best for them) and not much on the utility of the technologies as such. To this end, we make the following contributions:

- We conduct an empirical study on understanding patterns of influence on decision making when users are presented with different choices. The users observe this influence through a controlled information cascade in their social circle, i.e., their immediate neighbors in a social network.
- We find that compared to the control group where users were held to a single peer signal (reflecting one suboptimal technology) over all timesteps, users in a group who were treated early to large quantity of peer signals successfully opted for the suboptimal choice that their peers selected despite having their own knowledge about the utilities and the optimal choice. In fact when aggregated over the lifecycle of the game, we find this strategy of early substantial exposures to be far more effective than strategies which rely on delayed exposures to signals or situations where users observe uniformly increasing signals.
- Through multiple analyses on the effect of quantity of signals on user decisions, we conclude that unlike traditional studies which explain behavior diffusion as a factor of exposures ignoring the temporal patterns, the number of exposures alone does not explain successful social influence. Surprisingly, a sudden intervention in the form of stimulus through rise in quantity of signals momentarily succeeds in influencing more users to switch to the peer decision in subsequent rounds than strategies that uniformly increase signals over time.

The rest of the paper is organized as follows: we first discuss the related literature underpinning this study and the hypotheses that we will investigate in Section III. We present the experimental setup and methods designed for measuring social influence in Section IV.

## II. RELATED WORK

We are more often swayed by others’ decisions and behavior when we lack knowledge about the object of our decision, such as when we must choose a product that we do not know much about, is not well-described, or that we have little experience with. This is because the information we seek can be more cheaply acquired by observing others than by seeking it ourselves. Conventional studies suggest that as the consensus of entities in a social network increases – more signalers make the same signal – we assume that the information peers are conveying is valid and we are more likely to adopt the signaled behavior or decision [7].

**Number of signals:** The relationship between the number of signals an individual receives in its network, social influence, and the likelihood that said individual will adopt the behavior indicated by the signals is closely related to the linear threshold model in which an actor adopts a behavior after the signal count reaches an optimal threshold [11]. What, though, is the impact of repeated signals on the decision-making process, and more specifically how many signals are required to reliably influence an individual’s decision-making? There are mixed findings regarding the benefit of multiple exposures on the diffusion of information necessary to reach this threshold. People may prefer multiple confirmations from their peers to reassure themselves before making a decision [8], [9]. Experimental human studies using games from behavioral economics like the Prisoner’s Dilemma tend to find that the impact of zero to three signalers increases behavioral and decision-adoption in a linear fashion. However, debate still exists – some have found that repeated exposure to online signals in a social networking site slowed the subsequent spread of information [12].

**Information Parsing:** The manner in which an individual searches for information affects their decision-making. Individuals employ search strategies to reduce the number of choices [5]. This includes revising their initial opinions by processing and averaging the different influences acting on them [14] and social information provides one mechanism through which this is achieved. When social signals point towards a specific outcome or opinion, individuals will often adopt the opinions and behaviors of signalers, however while adopting behaviors based off social influence can be cost effective, it does not always lead to the most effective or efficient decision. Individuals must trade-off between trusting their own knowledge and trusting other’s opinions [15].

**Product Utility:** When making product decisions, outcomes related to the quality and need for a product change its utility or perceived value and therefore the risk of the selection. The need for a quality product also influences selection decisions. The perceived value of a product predicts how much an individual will search for information to inform their selection.

## III. THE PRESENT RESEARCH

We perform an online controlled experiment that manipulates the number of social signals and the signal pattern over time. We hypothesize that successful social influence

requires more than just receiving signals or exposures to information from peers, as both the utility of the technology and informational influences are at play. In our experiment, any decision made produces varying degrees of monetary gain based on utility. So, we speculate that successful social influence should be reflective of the mechanism through which information diffuses that ultimately instigates individuals to change their beliefs and therefore their adoption behavior. We present the following two hypotheses that we test in this paper with regards to the objective mentioned.

**HYPOTHESIS H1.** *Decision-makers will be more likely to choose cyber-defense providers without the optimal utility when they observe peers deviating from the optimal choice.*

This attempts to test the hypotheses studied before in [8] which show that behaviors spread to a larger portion of the population in a clustered network, indicating that additional social signals have significant effect on influence. However, the results on behavior diffusion reported by this paper are heavily associated with the clustered network organization more than the choice of the health behavior. Following this, in another study [18], authors show that the number of active neighbors is a positive indicator of influence, which is a similar finding reported by [8]. We test a richer version of these experiments by including the utility of the technologies and the monetary gains from different technologies as additional variables of interest to the decision making process conducted by individuals.

**HYPOTHESIS H2.** *The manner in which a decision-maker receives social signals or the pattern of influence will impact its decision when selecting the peer choice.*

Again, we note that the authors in [8] measure the effect of network structure in spread of influence, which differs from H2 we are investigating in this work, since we do not consider the position of the individual in the network. We consider that all peers of an individual are homogeneous with respect to the influence they can exert on it. The advantage of removing the network effects is that now we can control the diffusion of signals from the peers to an individual by administering them manually. This allows us to control the pattern of influence i.e. the number of signals sent over each time step to an individual.

We note that the basic component that distinguishes the two hypotheses is the focus on the selection of the *optimal* decision in H1 and the focus on the peer choice which we call the *influence* decision as would be described in details, in H2. As we point out later, a successful social influence constitutes situations where individuals not only deviate from the optimal decision but they also select the option that majority of their peers choose. We test H2 as a stronger version of H1 primarily to cater to this case.

#### IV. METHODS

To test our hypotheses, we ran an online, controlled decision-making game in which participants took on the role of a security officer at a bank. Participants were told that they and several of their peers at different banks were being asked to invest in a cyber-defense provider once a month for 18 months

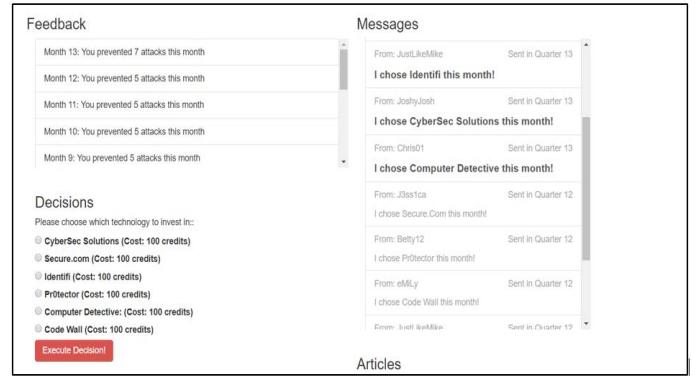


Fig. 2: Example screen in the cyber-defense provider selection task. Participants in the Uniform Messages (UM) condition of the study have access to a screen that looked like this. The *Feedback* section displays the number of attacks the participant prevented after each round. The *Decisions* section displays the 6 provider choices that it has. Finally, the *Messages* section is displayed after Round 12, from where on the participants can view what their peers selected in the previous round.

or rounds<sup>1</sup>. We separated participants into 5 groups based on pattern of social signal exposure which will be described in details in the *Design* subsection following this. For each group, we controlled the number of signals (corresponding to a suboptimal technology) that were sent to an individual from their peers over each time step. All participants could view brief descriptions of provider capabilities – e.g. one of them being “Secure.com utilizes algorithmic computer threat detection to keep systems safe. It prides itself on its efficiency and success rate in warding against attacks.” Participants were able to choose from 6 different providers - among which only one was optimal, preventing 7 attacks. The remaining 5 providers prevented 6 attacks each (from that perspective, all suboptimal technologies had the same utility). This information about the optimal and suboptimal providers is not available to the participants in any group. Furthermore, in all experimental conditions, they were able to see their peers’ decisions after 12 months or after Round 12. Figure 1 depicts the environment of the game, which was hosted by the Controlled Large Online Social Experimentation (CLOSE) platform and developed at Sandia National Laboratories [19].

The entire game was partitioned into two phases. For the first 12 rounds, no other information but a short excerpt about six potential providers was given. After the participants made their selection for a given month, they saw the number of attacks their provider had prevented in the corresponding period. For every attack they prevented, participants received \$0.02. Thus, they were incentivized to avoid more attacks and earn more money. However, since the participants have to explore the technologies to first acquire the knowledge of the technology utilities, the first 12 rounds allow for individual

<sup>1</sup>We use Rounds/Months/Timesteps interchangeably but which refer to one discrete unit of time in our study




Fig. 3: Illustration of the *linear cascade* diffusion. The technology  $C$  chosen by us as the suboptimal technology for influencing the user (in dots) cascades through the peers of the user over the 6 time steps. Colored nodes denote the activated peers w.r.t.  $C$  (manually preprogrammed by us) at each time step. Note that although at Rounds starting at 13 and ending at 18, there are subjects (uncolored) among peers who have not adopted  $C$ , their decisions (technologies adopted which may not be  $C$ ) at those time steps are visible to the subject in consideration (in dots). However, which users among the peers have been preprogrammed manually is by default unknown to the target subject.

decision making and exploration in the absence of any external knowledge about the technologies or peers.

In the second phase of the experiment which started at Round 13, we introduced social influence by allowing participants to see their peers' decisions after every round (see Figure 2), and by varying the pattern by which concentric decisions among the peers were made over time. We attempt to avoid network effects by using pre-programmed bots (these are the peers that the users see in their screen) and holding the network structure constant. All individuals in all groups have 6 neighbors whose choices could be viewed by the corresponding individuals. An example of the network structure is shown in Figure 3, where a participant receives social signals from its six neighbors about a random technology  $C$  among the 5 providers (barring the provider with the optimal technology). For each subject, this suboptimal technology  $C$  was selected as the peer choice that would be disproportionately signaled by peers over time (this pattern of influence would be manipulated by us). The motivation behind this deliberate selection of suboptimal  $C$  as the peer choice (controlled by us) is to investigate whether participants would be tempted to select the suboptimal choice. The knowledge of the peers' decisions allows a participant to rethink its own choices made in the presence of utilities and we use this to quantitatively measure the extent to which social influence is at play. We describe the signal patterns and conditions unique to each study in the following *Design* section.

For both the hypotheses  $H1$  and  $H2$ , the outcomes of interest are the decisions made by participants in the last six rounds, in the presence of social signals from peers. We explore whether decision-makers will be more likely to choose cyber-defense providers which are not the optimal choice when they have knowledge about the utilities and when they observe peers opting for choices which are not optimal. We note that people get feedback about their choice on the very next screen—and so choosing a technology during an intermediate timestep is more of a data-gathering exploration rather than

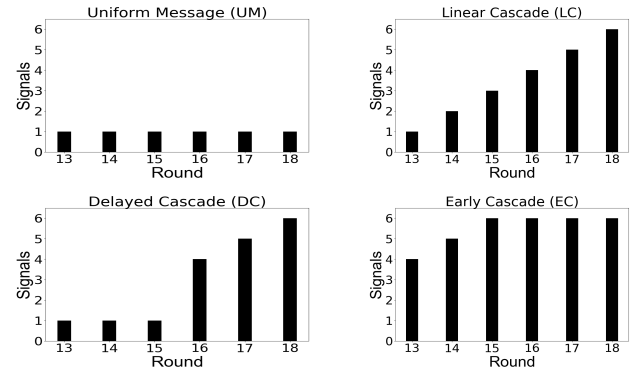


Fig. 4: The signal vs time step plots for the 4 patterns - note that for the NM group (not shown here), no peer signal in the form of pre-selected suboptimal technologies were sent to the participants at any time step.

their final choice. In order to allow for this initial bandwidth for exploration, we keep the first 12 timesteps (rounds) same for all subjects devoid of any interference. This helps in overcoming bias related to an individual's own knowledge about the utilities in the second phase of the experiment when they are treated to social signals.

#### A. Participants

We recruited a total of 357 participants for this study to play the same cyber-defense provider game. Participants were paid \$2 with the opportunity to earn up to \$4.52 since as mentioned before, they received a bonus of \$0.02 for every attack they prevented. Thus, the participants have a motivation to prevent more attacks in order to earn more money.

#### B. Design

Participants were randomly assigned to five groups with each group having unique members not involved in decision making as part of other groups. Decisions in the first 12 rounds (first phase) are made without social influence and are equal between groups. Let  $V(t)$  ( $t$  denotes a time step or round among the last 6 rounds) denote the number of peers of a user who at time  $t$  were programmed to select a chosen suboptimal technology by us. For the last six rounds participants in each group, except one group, receive signals in the following mechanism which we denote as the *patterns of influence* (Figure 4 shows the signal patterns for the groups):

- 1) *No Message (NM)*: Participants receive no message from the peers, so the last six round are exactly same for the participants as the first 12.
- 2) *Uniform Message (UM)*: Here we send only one signal using one peer (bot) of a participant at each round. So  $V = \{1, 1, 1, 1, 1, 1\}$  denotes uniform influence.
- 3) *Linear Cascade (LC)*: Here we incrementally activate one peer with the suboptimal technology at each step in round 2. So  $V = \{1, 2, 3, 4, 5, 6\}$  denotes uniformly increasing influence as shown in Figure 3.
- 4) *Delayed Cascade (DC)*: Here we send only one signal for the first 3 timesteps and send 4, 5 and 6 signals at



the last 3 timesteps respectively. So  $V = \{1, 1, 1, 4, 5, 6\}$ . The objective is to see whether the sudden change in the number of signals acts as a catalyst for increasing successful influence at the later stages of the experiment.

- 5) *Early Cascade (EC)*: In this setup, we send more signals in the beginning. So  $V = \{4, 5, 6, 6, 6, 6\}$ . This pattern allows us to ask if an early trigger is able to sustain the levels of influence, or whether participants will return to the optimal choice at the later stages.

Note that in all conditions, users can switch back to any choice in the next round after having selected an option in the current round. We consider the NM and UM groups as our baseline groups and LC, EC, DC groups as our treatments groups of interest. Note that while the pattern remains the same, the suboptimal technology chosen by us for each subject and that cascades through its peers is random. An example of the Linear Cascade (LC) pattern is shown in Figure 3, where the subject (marked in dots) is able to receive social signals from its 6 neighbors. At Round 13 (start of the second phase), a signaler (node v1) selects a suboptimal provider C, and over the next five rounds, the remaining peers adopt the same behavior one after another (we will refer to this as the influence decision). Note that although we program only a few of the peers (bots) to adopt  $C$  over time, the subjects are able to view all the bots and their decisions in their dashboard for all the last 6 timesteps. We emphasize that all the peers of the subjects are bots and do not share any topology between themselves, thereby we sideline the effects of network on the individual behavior changes.

## V. ANALYSIS

| Group | # participants | Average number of attacks prevented |
|-------|----------------|-------------------------------------|
| NM    | 55             | 105.2                               |
| UM    | 71             | 106.28                              |
| LC    | 79             | 103.81                              |
| DC    | 81             | 104.8                               |
| EC    | 71             | 103.83                              |

TABLE I: Average number of attacks prevented by subjects in each group. The lower attacks suggest participants deviated more from the optimal decision responding to social influence.

### A. Distribution of attacks prevented

Table I shows the distribution of attacks prevented by subjects in each group. We observe that, on average subjects in the EC and LC groups prevent more attacks compared to others. Based on a survey analysis, we found that none of the traits like computer anxiety, computer confidence, computer liking, intuition or neuroticism were correlated to the number of attacks prevented in all groups.

### B. Distributions of decisions by individuals

As a first step towards investigating hypothesis  $H1$ , we analyze the kinds of social signals or the cyber security




Fig. 5: Decisions chosen by the bots for participants in each group (here groups are considered separately). Note that only the decisions that are not optimal are sent as social signals in the second phase of the experiment.

technologies (which were not the optimal technology) chosen by the bots of each participant and whether they are uniform across all the groups. In our experimental setup throughout, technology provider 1 or decision 1 constituted the optimal choice preventing 7 attacks, while the rest of the technologies prevented 5 attacks and are being termed as suboptimal choices. From Figure 5, we observe that the random selection of a suboptimal technology  $C$  for a participant (and which differs for each participant) introduces some disproportionate selection of the 5 possible suboptimal technologies, when considering individual groups. For UM group, around 25% of users received decision 4 (maximum among all suboptimal technologies) while 29% of users in the early cascade group were sent decision 5 (maximum) as their peer choices and 28% of users in the delayed group cascade were sent decision 6 (maximum). However, we see that users in the linear cascade group had signals sent that are uniformly distributed among the population.

Having observed that there was not one pre-programmed peer choice as the strategical obvious suboptimal choice across all groups, we proceed with investigating  $H1$ . In order to detect any implicit occurrence of a selection bias over the participants, we analyze whether there is any significant difference in the groups with respect to choices made in the first 12 rounds. To accomplish that, we plot the probability that an individual makes each decision when aggregated over the first 12 rounds. We find that there is clearly no evidence of differences in the mean statistics of the probability distributions between the treatments groups (LC, DC, EC) and the control groups (UM, NM) (Refer to Figure 1 in the Appendix<sup>2</sup>). The results (p-values for each group) in Tables 1, 2 and 3 in the Appendix suggest no significant difference in the distributions among the groups. This rules out any bias among the participants themselves in the absence of externalities. However, in the second phase of the experiment (Rounds 13 to 18 aggregated), we find differences in the selection patterns among the decision-makers in their respective groups. We find the following observations from Figures 6(a) and (b) for our treatment groups (see Tables in the Appendix):

- 1) **LC**: With respect to the NM group, there are no sta-

<sup>2</sup>Online **Appendix** can be accessed here: [Link](#)

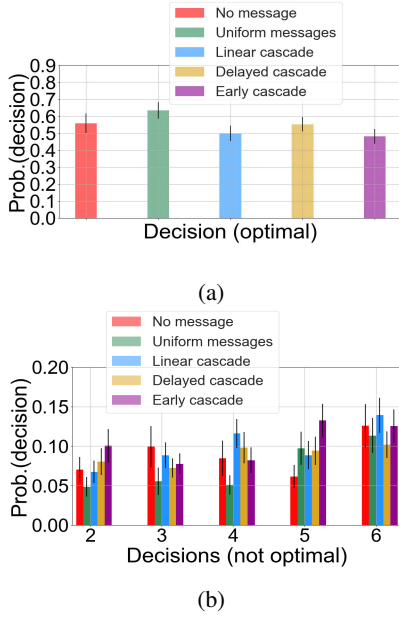


Fig. 6: Probability of decisions made in Rounds 13 to 18. (a) Probability of making the optimal decision. (b) Probability of making the suboptimal (other 5) decision. The error bars denote standard error over the distributions.

tistically significant differences in the decisions taken by the participants in LC - we carried out a similar statistical test comparing the group pairwise means as done for the first 12 steps. On the other hand, we find that there is a statistically significant difference in the means of the probabilities compared to the UM group participants for the optimal decision ( $p=0.04$ ,  $df=149$ ) at  $\alpha = 0.1$ . The difference shows that a significantly less number of participants are tempted to choose the optimal technology provider in the presence of linear cascading signal pattern than when a single signal is sent across all timesteps.

- 2) **DC:** For participants in the DC group we find that w.r.t. the NM group, there is a statistically significant difference for the participants who chose decision 5 ( $p=0.09$ ,  $df=140$ ) and with respect to UM group, the participants opting for decision 2 and 4 differed significantly from those in DC. However, w.r.t NM group or UM group participants, we observe that the users do not differ in their selections when it comes to choosing the optimal provider. Also it does not differ with respect to the most common choice among the bots for DC group - the decision 6 shown in Figure 5.
- 3) **EC:** With respect to the UM group users, the users in this group differ in the choice most selected by bots - decision 5 ( $p=0.07$ ,  $df=128$ ). This is a successful case of social influence when considering the macro-adoption process for the group as the users not only steered away from the optimal choice, but they also steered towards the decision of their peers.

This suggests that while the social signal to some extent




Fig. 7: Fraction of users in each group adopting the *influence* decision chosen by their peers.

influences an individual in the LC group to deviate from the optimal choice, it does not always translate to the peer decision that was chosen for the corresponding individual. It rather gears the user towards more exploration. However, an early burst of signals in the EC successfully translates towards social influence wherein the users sway towards the influence decision more.

### C. Degree of Influence

This first analysis of H1 does not shed any light on the temporal variations in the decision making process exhibited by users in different groups - it shows some statistically significant differences in the choices made for specific decisions (or technology providers) over the entire second phase. While it did show that not all cascade patterns successfully influenced users towards not choosing the optimal provider, leading them to the more popular bot choices, it brings up the question that is posed for hypothesis H2: what constitutes successful influence and if so, does the manner in which the signals are sent determine successful influence?

To this end, we analyze how the proportion of users who switch to the influence decision, ignoring the optimal choice changes, throughout the rounds for each group. We note that for each user the *influence decision* is randomly selected before the second phase starts. Figure 7 shows the fraction of users in each group adopting the influence decision from rounds 13 to 18, when the experimental participants observe their peers' decisions. Following from the observations regarding H1, we find that the probability of successful influence for participants in EC has the strongest effects on decision-making. Participants in EC are most likely to deviate from selecting the optimal provider as shown in Figure 6. We also find that, while for the first 3 timesteps in the second phase, participants in EC group exhibit the maximum adoption compared to other groups, Round 15 had the maximum retention, where participants were exposed to all their peers selecting the same technology. - this may be due to new users adopting the influence decision or due to the same users from previous rounds who do not switch back. We will explore this in the following sections.

The participants in the DC group exhibit successful response towards the sudden increase in signals at Round 16 which is shown by a 65% increase in user adoption of the

*influencedecision* compared to the previous step. These results become close to the 25% of users making peer decision selection at Round 16 and which is also the maximum among all groups surpassing the early cascade adoption ratio. However, there is no substantial increase in the adoption fraction for the users in the linear cascade group - we do note that the adoption peaks at Round 15 for the linear cascade users before it drops again. These observations from Figure 6 and Figure 7 suggest that while an early burst of social signals successfully persuades users in EC to adopt the influence decision, causing an aggregated overall maximum selection of the peer decision, a sudden impulse in the quantity of social influence also successfully steers users towards the influence decision in the later stages. As a side experiment to measure the degree of drift away from optimal decision, we also analyze the fraction of users who shift away from decision 1 (the optimal) at each timestep similar to what has been shown in Figure 7. However, we do not find any clear distinctions among the groups in terms of the fraction of subjects who move away from optimal decision. This leads us to conclude that our experimental setup had a more pronounced effect when influencing users towards their peer social signal or one particular suboptimal choice than just influencing them away from the optimal decision - the fact that the users eventually move away from the optimal decision does not contribute much in distinguishing the patterns of influence.

#### D. Measuring the effect of quantity of signals

In this section, we investigate whether the quantity of signals alone stand out as the sole factor of influence or does the manner in which the signals are sent, contribute more to the social influence. Before going into the analysis, we define a few notations: we denote a subject in this study as  $u$  where  $u$  can belong to any group. We denote the decision taken by a subject  $u$  at time  $t$ ,  $t \in [1, 18]$  by  $d_u(t)$ . Let the *influence decision* (the suboptimal technology  $C$ ) selected by the peers of  $u$  be denoted by  $C_u$  (note again  $C_u$  would be different for each participant  $u$ ). We define the number of peers of  $u$  who reflect the influence decision  $C_u$  at round  $t$  as  $V_u(t)$  ( $V_u(t)$  would be different for users in each group, for e.g. for a  $u$  in LC group,  $V_u(t = 1) = 1$ ,  $V_u(t = 2) = 2$  and so on). We define  $T_{treat}$  as the sequence of time points during which the subjects receive signals from their peers i.e.  $T_{treat} = [13, 18]$ . Following this, we denote the timestep  $t$ ,  $t \in T_{treat}$  when an individual  $u$  first switches to  $C_u$  as  $t_u^f$  (by definition,  $t_u^f \in T_{treat}$ ). For each signal quantity  $s$ , we measure the proportion of individuals (in each group) who made their **first** switch to their peer decision only after they were exposed to  $s$  signals. Formally, for any  $t \in T_{treat}$ ,  $R(s) = \frac{|\{u \mid d_u(t)=C_u \wedge V_u(t)=s \wedge t=t_u^f\}|}{|\{u\}|}$  (R(1) in LC group denotes the number of individuals in LC group who made their first switch to their peer influence decision after being exposed to just 1 signal. Similarly,  $R(6)$  in EC group denotes the number of individuals in EC group who made their first switch to their peer influence decision only after being exposed to 6 signals, so this can happen in any

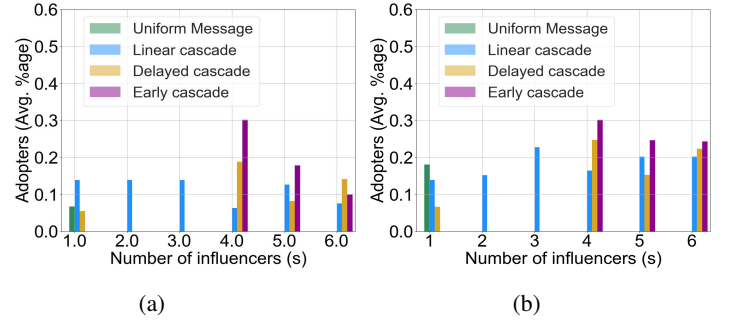


Fig. 8: Plots of adoption under the influence of  $s$  exposures. (a) For each signal quantity  $s$ , proportion of individuals who made their **first** switch to their peer decision after being exposed to  $s$  signals, (b) Proportion of individuals who adopt their peer decision after being exposed to signals from  $s$  influencers.

of the last 4 rounds in EC). The denominator in the formula here denotes the number of individuals in the group. We bin the values from  $R(s)$  based on  $s$  and take the mean for each group, since for some groups, there can be multiple rounds with the same number of exposures or signals. From Figure 8(a), we observe that for EC group,  $\sim 30\%$  of users under the influence of 4 signalers, for DC,  $\sim 18\%$  of users at 4 signalers and for LC,  $\sim 15\%$  of users at 3 signalers (all these being the maximum ratio) made their **first** switch to the peer decisions in the second phase of the experiment. However, on close observation, we find that the number of exposures alone does not explain the adoption behavior. When we compare the EC participants with those in DC group, we find that with 4 exposures in  $T_{treat}$  (at round 13 for EC and at round 16 for DC), the proportion of adopters in EC making their first peer decision switch (30% of users) is higher than the proportion in DC (18% of users), although the same 4 quantity of signals are delivered at different timesteps for the 2 groups. However, while there is a constant decrease in the number of adopters making first switches in the EC group going from 4 to 6 exposures, we see that the DC influence pattern does not decrease the same way for 4 to 5 to 6 exposures. This suggests that the sudden stimulus from the delayed exposure somehow succeeds in influencing more users to make their first switch to the peer decision compared to the EC group (note that all the users under these different quantities are unique since we measure their first switch). On the other hand, for the linear cascade group we do not find one quantity that is most effective in the influence - in the LC setup, there is no one exposure that impacts the adoption behavior the most in terms of successful influence.

In addition, we measure the cumulative adoption ratio for each group defined as: for any  $t \in T_{treat}$ ,  $Z(s) = \frac{|\{u \mid d_u(t)=C_u \wedge V_u(t)=s\}|}{|\{u\}|}$ . In simple terms, it measures the number of individuals in a group who adopt the influence decision under  $s$  exposures irrespective of whether it is the first switch. This is demonstrated in Figure 8(b). When we combine the results obtained in Figure 8(a) with Figure 8(b),






Fig. 9: Success Ratio.

we find an interesting observation for the LC group. The linear cascade pattern is able to retain most of the users even after first switch at timestep 13 as the cumulative ratio increases up to 3 exposures (which occurs at timestep 15). This suggests that the LC pattern is effective in terms of retention in the early stages of the cascade.

#### E. Dynamics of adoption

We end this study with a retrospective analysis to understand the dynamics of influence. In an attempt to quantify the effect of the influence on subjects in a more constrained setting, we measure at every timestep, the ratio of individuals who adopted the peer decision at the respective timestep to the number of individuals (in their respective group to which it belongs) who switched to their influence decision *at least once* within their lifecycle ( $T_{treat}$ ). Note that this is different from previous measures in 2 ways: first we retrospectively filter out users who never adopted their peer decision (in the real world these are users who would not be susceptible to influence or are immune as such). Second, we analyze this ratio at the end of their exploration phase, in Round 18, when everybody have supposedly settled down. We define a symbol  $N(u)$  as the number of time steps for which a user  $u$  adopts  $C_u$  in  $T_{treat}$  (this is measured retrospectively aggregating all timesteps beforehand). Formally it is defined as:  $Success\ ratio(t) = \frac{|\{u \mid d_u(t)=C_u\}|}{|\{u \mid N(u) \geq 1\}|}$ . The denominator denotes the number of individuals who have adopted the peer decision at least once from Rounds 13 to 18. The comparison shown in Figure 9 among the four groups (the No message group does not have any influence decision) demonstrates that while the EC group adopts the influence decision more quickly than other groups, the stimulus in signals quantity at Round 16 in DC group affected the participants. This is confirmed when the effects of DC strongly outstrip those observed from EC in the last round where both groups receive 6 signals.

#### VI. CONCLUSION

We present a controlled experiment that demonstrates how individuals can deviate from the optimal choices as a result of social influence. While an early cascade influences decision-makers to deviate most from their choice when aggregated over the game, the speed with which individuals switch to their peers' decision further increases when an impulsive stimuli is present. Such conclusions can have diverse impacts on real world use-cases where we can devise strategies to influence

people towards making a better choice even when there is little motivation.

**Acknowledgments.** Some of the authors are supported through the ARO grant W911NF-15-1-0282. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. We would like to thank Glory Emmanuel-Avina and Victoria Newton for their insights and assistance with preliminary experiment design.

#### REFERENCES

- [1] Graf-Vlachy L, Buhtz K, König A. Social influence in technology adoption: taking stock and moving forward. *Management Review Quarterly*. 2018;68(1):37–76.
- [2] Bakshy, Eytan, et al. "Everyone's an influencer: quantifying influence on twitter." *Proceedings of the fourth ACM international conference on Web search and data mining*. ACM, 2011.
- [3] Kraut RE, Rice RE, Cool C, Fish RS. Varieties of social influence: The role of utility and norms in the success of a new communication medium. *Organization Science*. 1998;9(4):437–453.
- [4] Graham, D. "Experimental studies of social influence in simple judgment situations." *The Journal of Social Psychology* 56.2 (1962): 245-269.
- [5] Payne JW, Braustein ML. Risky choice: An examination of information acquisition behavior. *Memory & Cognition*. 1978;6(5):554–561.
- [6] Romero, Daniel M., et al. "Influence and passivity in social media." *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, Berlin, Heidelberg, 2011.
- [7] Bearden WO, Netemeyer RG, Teel JE. Measurement of Consumer Susceptibility to Interpersonal Influence. *Journal of Consumer Research*. 1989;15(4):473–481.
- [8] Centola D. The spread of behavior in an online social network experiment. *science*. 2010;329(5996):1194–1197.
- [9] Zhou C, Zhao Q, Lu W. Impact of repeated exposures on information spreading in social networks. *PloS one*. 2015;10(10):e0140556.
- [10] Dickinger A, Arami M, Meyer D. The role of perceived enjoyment and social norm in the adoption of technology with network externalities. *European Journal of Information Systems*. 2008;17(1):4–11.
- [11] Myers, Seth A., Chenguang Zhu, and Jure Leskovec. "Information diffusion and external influence in networks." *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2012.
- [12] Ver Steeg G, Ghosh R, Lerman K. What stops social epidemics? In: *ICWSM*; 2011.
- [13] Anagnostopoulos, Aris, Ravi Kumar, and Mohammad Mahdian. "Influence and correlation in social networks." *ACM KDD* 2008.
- [14] Muchnik L, Aral S, Taylor SJ. Social influence bias: A randomized experiment. *Science*. 2013;341(6146):647–651.
- [15] Schöbel M, Rieskamp J, Huber R. Social influences in sequential decision making. *PloS one*. 2016;11(1):e0146536.
- [16] Bao P, Shen HW, Chen W, Cheng XQ. Cumulative effect in information diffusion: empirical study on a microblogging network. *PloS one*. 2013;8(10):e76027.
- [17] Grujić J, Eke B, Cabrales A, Cuesta JA, Sánchez A. Three is a crowd in iterated prisoner's dilemmas: experimental evidence on reciprocal behavior. *Scientific reports*. 2012;2:638.
- [18] Zhang J, Liu B, Tang J, Chen T, Li J. Social Influence Locality for Modeling Retweeting Behaviors. In: *IJCAI*. vol. 13; 2013. p. 2761–2767.
- [19] Lakkaraju K, Medina B, Rogers AN, Trumbo DM, Speed A, McClain JT. The Controlled, Large Online Social Experimentation Platform (CLOSE). In: *International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction*. Springer; 2015. p. 339–344..