Influence of Controlled and Uncontrolled Interventions on Twitter in Different Target Groups

Peter-Paul van Maanen*†, Olav Aarts*, Erik Boertjes* and Remco Wijn*
*Netherlands Organisation for Applied Scientific Research (TNO), The Netherlands
Email: {peter-paul.vanmaanen,olav.aarts,erik.boertjes,remco.wijn}@tno.nl
†Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

***Abstract*—In this paper the influence of interventions on Twitter users is studied. We define influence in a) number of participants, b) size of the audience, c) amount of activity, and d) reach. Influence is studied for four different target groups: a) politicians, b) journalists, c) employees and d) the general public. Furthermore, two types of interventions are studied: a) by all Twitter users (i.e., uncontrolled interventions), and b) those tweeted by an organization that benefits from any resulting influence (i.e., controlled interventions). As a case study, tweets about a large Dutch governmental organization are used.**

Results show a clear relation between the number of uncontrolled interventions and influence in all four target groups, for each of the defined types of influence. Controlled interventions show less influence: Significant influence was found for the general public, but influence for politicians and employees was only mildly significant, and no influence was found for journalists. The effect found for uncontrolled interventions however suggests that this influence is indeed reachable for some target groups, even when the number of interventions is small, and very well reachable for all target groups, provided the number of interventions is large enough.

In addition to this we found that interventions influence groups to a different extent. Own employees were influenced strongest, differing significantly from the other groups.

***Keywords*-Social Media; Twitter; Social Influence**

I. Introduction

Image is important to institutions and companies. How the public views an institution strongly affects support for this institution. For this reason, many companies are motivated to monitor and influence public opinion. With the rise of social networks such as Twitter new methods for monitoring and influencing the public's opinion are gaining popularity. This paper contributes to existing research and practice on information diffusion through social networks by introducing new influence measures for Twitter that we applied to measure the influence of tweets sent out by a large Dutch governmental organization. We aim to elucidate how and to what extent organizations can influence users in their online social network and differentiate different types of users based on their professional affiliation.

II. Background

What constitutes influence on social media is focus of debate by both researchers and practitioners. Huberman, Romero and Wu distinguish between friends and followers [1]. According to these authors, a friend is another user to which the user has directed two or more tweets (using the "@username" convention). Comparisons show that users' interactions with friends reflect a different network than following relationships suggest [2]. This leads Yang and Counts to conclude that the interaction network, rather than the followers network, is preferable for network analyses of Twitter [3]. This idea is echoed by practitioners [4] and seems to become a truism among those working on social network analyses.

However, a study by Mehta, Mehta, Chheda, Shah and Chawan [5] contradicts this idea. Using sentiment analysis, these researchers conclude that the number of followers of a user is a better predictor than the number of retweets and lists to estimate the influence of a message. Other research shows that although stronger ties (e.g., friends) are individually more influential, weak ties (e.g., followers) are responsible for the propagation of novel information [6]. This suggests that weak ties may play a more dominant role in the dissemination of information online than currently believed [7].

Based on the above findings, we suggest that what constitutes effective influence on social media depends on the specific goals of the communication. On the one hand, given that individuals influence each other more the closer their interpersonal bonds are, individuals will be persuaded easier by friends than by mere followers [8], [9]. Thus, when the goal is to persuade, influence measures based on the friend network (vs. the follower network) should lead to best predictions of influence. On the other hand, when the goal is to inform, source credibility is less of an issue [10], [11] and the weight shifts to the reach of a tweet which may be determined best using the followers network.

In the present paper the above two views are translated into four types of influence (see Table I). First, elaborating on [1], we define *participants* as the part of the friends network that is actively participating in a thread. In other words, participants are the users in the friends network filtered on the topic of interest (e.g., #topic) within a specified time interval. Second, elaborating on [5], we define *audience* as the part of the follower network that is passively following the participants. That is, when user B and user C react to user A, then followers of user A, user B and user C are counted as part of the audience.

Parallel to these two influence types that focus on the number of users (i.e., a part of the *network*), we suggest that influence is also related to the number of tweets these users

TABLE I
DIFFERENT INFLUENCE TYPES.

	active	*passive*
network	participants	audience
messages	activity	reach

sent out (i.e., the number of *messages*). To count as a friend or follower in a discussion requires a user to have sent at least one tweet in response to someone else, but it is very well possible that individuals respond with more than one tweet and this may have a different influence on friends and followers [12], [13]. For this reason, as a third influence type we suggest the total number of messages which are actively sent by participants. In Table I this is referred to as the *activity*, which is the number of messages sent in the discussion. Finally, as a fourth influence type we suggest the relation between activity and the followers network. In Table I this is referred to as the *reach*, which is the part of the followers network that is passively following the activity. That is, the number of times each message, as part of the discussion, is followed.

III. HYPOTHESES

For organizations it may be of interest to know more about how measures of influence apply to different target groups. For example, organizations may wish to influence public opinion to generate traction for a campaign, or governmental organizations may try to prevent budget cuts by influencing the wider public or journalists as they are likely to influence politics [14], [15]. In effect, in such situations the organization will likely be interested in knowing to what extent they have been able to reach their goals of influencing these parties. In this paper we focus on governmental organizations and their influence on online social media. For this reason, we distinguish between influence on the following target groups: *politicians*, *journalists*, the organization's own *employees*, and the *general public* (i.e., people not in one of the other categories).

Because a basic feature of social networks is that content is shared among its members, often with the goal of informing, influencing, or eliciting reactions from others, we distinguish between primary tweets and reactions to these tweets. In this paper we regard primary tweets on the organization in a discussion thread as *interventions* (that is, we regard these tweets as attempts to inform, influence or elicit reactions from the Twitter network). We further distinguish between *controlled* interventions (i.e., interventions sent out by the organization itself) and *uncontrolled* interventions (i.e., interventions sent out by Twitter users other than the organization).

Based on the above, we hypothesize that:

Hypothesis 1. *Controlled and uncontrolled interventions have a positive effect on our four influence types in each of our four target groups.*

In public opinion formation it makes sense that not everyone is equally receptive to interventions. Generally, people related to or interested in the person or organization making an intervention, or individuals with a strong interest in the activity displayed by this person or organization, will be influenced more strongly. We therefore assume that the effect of interventions on influence measures is moderated by type of group. We take the general public as a reference to compare to other groups.

Specifically, we hypothesize that:

Hypothesis 2. *(Un)controlled interventions have stronger influence on the organization's own employees than on the general public.*

Further, we assume that politicians have party programs to adhere to, making them potentially reluctant to get involved in controversial topics on social media.

We therefore argue that:

Hypothesis 3. *(Un)controlled interventions have weaker influence on politicians than on the general public.*

As part of their functions, journalists are generally looking out for news and in-depth information on a variety of topics. We therefore expect that journalists will show more interest in the functioning and current state of affairs of the organization than the general public.

We therefore argue that:

Hypothesis 4. *(Un)controlled interventions have stronger influence on journalists than on the general public.*

IV. METHOD

A. Use case

For this study we used the case of a large Dutch organization, anonymously called "Dutchorg". Dutchorg is an important part of Dutch society, but lacks the attention it wants to have in Dutch politics. To prevent any further lack of attention, Dutchorg wants to investigate the possibilities of intervening through social media (in this case Twitter). These interventions are aimed to increase positive influence on different target groups, such as politicians.

B. Data gathering

We used the Twitter stream API to harvest all tweets with #dutchorg during a two week period, starting one week before, and ending one week after the Dutch elections in 2012. Since the election issues were closely related to #dutchorg, we expected the majority of tweets in that period.

We gathered 7146 tweets, sent by 1270 distinct users. One of these users sent 14 tweets that were specifically meant to alter the public opinion about Dutchorg. These (controlled) interventions were done twice a day, over a period of one week.

The streaming API provided, for each tweet, its text, the account name of the user that sent the tweet, the time of sending, the location (if available), and the number of followers and followees of the user.

C. Independent variables

In this study we investigate the influence of interventions on Twitter for different *intervention types* and for different *target groups*.

1) Intervention type: Influence can be determined for *uncontrolled interventions* or for *controlled interventions*.

a) Uncontrolled interventions: These interventions are tweets that contain "#dutchorg" and are original. Tweets are considered *original* when they are not retweets (tweet does not include "RT:") or mentions (tweet does not include "@"). The *intervention scale* is determined by the number of such tweets within each time interval.

The time interval used in this study (i.e., between $t-1$ and t) is *one hour*, since this is considered a small enough interval to detect any changes in influence, and large enough to have enough data to say anything about influence.

The following is used to calculate the intervention scale:

$$\text{u\_interventions}(t) = |I(t-1,t)|$$

where $I(t_1, t_2)$ is the set of original tweets between time point t_1 and t_2.

Note that a message cannot be counted as an intervention during multiple time intervals, and that therefore it holds that:

$$\sum_{t=1}^{t_e} \text{u\_interventions}(t) = |I(0,t_e)|$$

where t_e is the last time point in the data.

b) Controlled interventions: These interventions are those tweets that are deliberately made by Dutchorg itself to increase influence on different target groups. They are a (much) smaller subset of the set of uncontrolled interventions.

2) Target group: A *target group* is defined as a portion of the total set of Twitter users that belongs to a certain group of people in real life. In this study these groups are *politicians*, *journalists*, *employees* and *general public*. For this, the set of Twitter users was annotated by hand. Both individual Twitter accounts as well as other online sources (such as LinkedIn or personal webpages) were used to determine to which group each user belongs. A user can belong to more than one group, but not to no group at all. When there was any doubt, a user was included in the general public group.

D. Dependent variables

1) Influence: In line with Table I, influence of Twitter interventions was measured in four ways: by means of the *participants*, *audience*, *activity* and *reach*.

Before measuring influence, the following two steps have been made to prepare the data for analysis (from here on called *influence data*):

1) All data was separated from the data used for the calculation of the independent variable.
2) The data was filtered on the used target group.

The first step was to make sure the independent and dependent variables are not in principle already related to each other (i.e., number of interventions increase and therefore possibly a large part of the number of tweets that determine influence also increase).

For uncontrolled interventions, the second step resulted in a target group size for employees of 287, for journalists of 41, for politicians of 140 and for the general public of 573 (for controlled: 308, 58, 151 and 726, respectively).

a) Participants: Participants are the part of the Twitter network that is (actively) tweeting about Dutchorg. The number of participants is determined by the number of users in the influence data within each time interval (one hour).

The following is used to calculate the influence on the basis of participants:

$$\text{u\_influence}_{P_g}(t) = |P_g(t-1,t)| \cdot \frac{|P(0,t_e)|}{|P_g(0,t_e)|}$$

where $P(t_1, t_2)$ is the set of participants between time point t_1 and t_2, $P_g(t_1, t_2)$ those in target group $g \in \{\text{p}, \text{j}, \text{e}, \text{g}\}$, and t_e is the last time point in the data.

The factor $\frac{|P(0,t_e)|}{|P_g(0,t_e)|}$ is a normalization constant that makes it possible to compare influences between target groups on each time point t. The influence on a specific target group g is normalized to a value *as if* the total number of participants of g (that is, $|P_g(0,t_e)|$) is equal to that of all participants (that is, $|P(0,t_e)|$). This is also done in the calculations of the subsequent influence types.

Note that participants can participate during multiple time intervals, and that therefore it holds that:

$$\begin{aligned}\sum_{t=1}^{t_e} \text{u\_influence}_{P_g}(t) &= (|P_g(0,1)| + |P_g(1,2)| + \ldots \\ &\quad |P_g(t_e-1,t_e)|) \cdot \frac{|P(0,t_e)|}{|P_g(0,t_e)|} \\ &\geq |P_g(0,t_e)| \cdot \frac{|P(0,t_e)|}{|P_g(0,t_e)|} = |P(0,t_e)|\end{aligned}$$

b) Audience: The audience is the part of the Twitter network that is (at least passively) following Dutchorg. The audience is determined by the sum of all followers of each participant within each time interval (one hour).

The following is used to calculate the influence on the basis of the audience:

$$\text{u\_influence}_{A_g}(t) = \sum_{p \in P_g(t-1,t)} |\text{followers}(p)| \cdot \frac{\sum_{p \in P(0,t_e)} |\text{followers}(p)|}{\sum_{p \in P_g(0,t_e)} |\text{followers}(p)|}$$

where $\text{followers}(p)$ is the set of followers of participant p.

c) Activity: The activity is the scale of the messages about Dutchorg. Activity is determined by the total number of tweets in the influence data within each time interval (one hour).

The following is used to calculate the influence on the basis of activity:

$$\text{u\_influence}_{M_g}(t) = |M_g(t-1,t)| \cdot \frac{|M(0,t_e)|}{|M_g(0,t_e)|}$$

where $M(t_1, t_2)$ is the set of messages between time point t_1 and t_2, $M_g(t_1, t_2)$ those in target group g, and t_e is the last time point in the data.

Note that a message cannot be counted as activity during multiple time intervals, and that therefore it holds that:

$$\sum_{t=1}^{t_e} \text{u\_influence}_{M_g}(t) = |M(0,t_e)|$$

d) Reach: The reach is determined by the sum of followers of each actor of each tweet in the influence data within each time interval (one hour).

The following is used to calculate the influence on the basis of reach:

$$\text{u\_influence}_{R_g}(t) = \sum_{m \in M_g(t-1,t)} |\text{followers}(m)| \cdot \frac{\sum_{m \in M(0,t_e)} |\text{followers}(m)|}{\sum_{m \in M_g(0,t_e)} |\text{followers}(m)|}$$

where followers(m) is the set of followers of message m.

E. Statistical data analyses

For each intervention type and each target group (independent variables) we measured each of the four influence variants (dependent variable).

1) Uncontrolled interventions: First we analyzed the influence of uncontrolled interventions in all four target groups. This was done by linear regression of intervention scale and influence. This means that a line was fitted in the scatter plot of pairs $(\text{u\_interventions}(t), \text{u\_influence}_{i_g}(t))$, for each influence type i and target group g. This was used to test Hyp. 1 for uncontrolled interventions. As Hyp. 1 predicts "some" positive influence of interventions on each target group (as opposed to a specific positive influence), linear regression was used here (as opposed to regressing polynomials with degrees higher than one).

Then, within each influence type i, the possibly statistically significant regression lines were compared to each other to see if there are any significant differences. This was used to test Hyp. 1, 2, 3 and 4 for uncontrolled interventions.

2) Controlled interventions: Secondly, we analyzed the influence of controlled interventions. This was done by averaging the differences between the uncontrolled influences an hour before and after each controlled intervention:

$$\sum_{t_c \in T_c} \frac{\text{c\_influence}_{i_g}(t_c)}{|T_c|} = \sum_{t_c \in T_c} \frac{\Delta\text{u\_influence}_{i_g}(t_c)}{|T_c|} = \sum_{t_c \in T_c} \frac{\text{u\_influence}_{i_g}(t_c+1) - \text{u\_influence}_{i_g}(t_c)}{|T_c|}$$

for influence type i, target group g and controlled intervention time point t_c. These averages were then compared to 0 to see if there indeed was influence of the controlled intervention on a specific target group. This was used to test Hyp. 1 for controlled interventions. Note that, as described earlier, different influence data were used for controlled interventions.

Also, for the controlled interventions, within each influence type i, the influences between the target groups were compared to each other to see if there are any significant differences. Again, this was used to test Hyp. 1, 2, 3 and 4, but then for controlled interventions.

V. Results

A. Uncontrolled interventions

1) Relating interventions to influence: Fig. 1 shows the twenty regression lines after regressing the four influence types on intervention scale for the four target groups.

As shown in Table II the regression results (with $n = 304$ and $df = 303$) show a significant relation between the intervention scale and influence (all four types) for all four target groups. Hence, uncontrolled interventions have influence on all target groups. Hyp. 1 could therefore be accepted for uncontrolled interventions, for each influence type and target group.

2) Comparing influences per target group: Four separate one-way analyses of variance (ANOVA) showed significant main effects of target group, for participants, $F(3,1) = 33.26$, $p = 0$ and audience, $F(3,1) = 10.01$, $p = 0$, but not for activity, $F(3,1) = 0.2$, $p = .90$, and reach, $F(3,1) = 0.42$, $p = .74$. Tukey's honestly significant difference post hoc analyses were therefore only conducted for the participants and audience influence types.

a) Participants: For the participants, uncontrolled interventions had more influence on employees than on the general public. Hyp. 2 could therefore be accepted for participants. Influence on politicians was not weaker than on the general public (no difference and weaker than on employees). Influence on journalists was not stronger than on the general public (no difference and weaker than on employees). Both Hyp. 3 and 4 could therefore not be accepted for uncontrolled interventions and participants.

b) Audience: For the audience, uncontrolled interventions had more influence on employees than on the general public. Hyp. 2 could therefore be accepted for audience. Influence on politicians was not weaker than on the general public (no difference and stronger than on journalists). Influence on journalists was not stronger than on the general public (no difference and weaker than on journalists and employees). Both Hyp. 3 and 4 could therefore not be accepted for uncontrolled interventions and audience.

B. Controlled interventions

1) Relating interventions to influence: In Fig. 2 the effect of controlled interventions on all four influence types for all four target groups is shown.

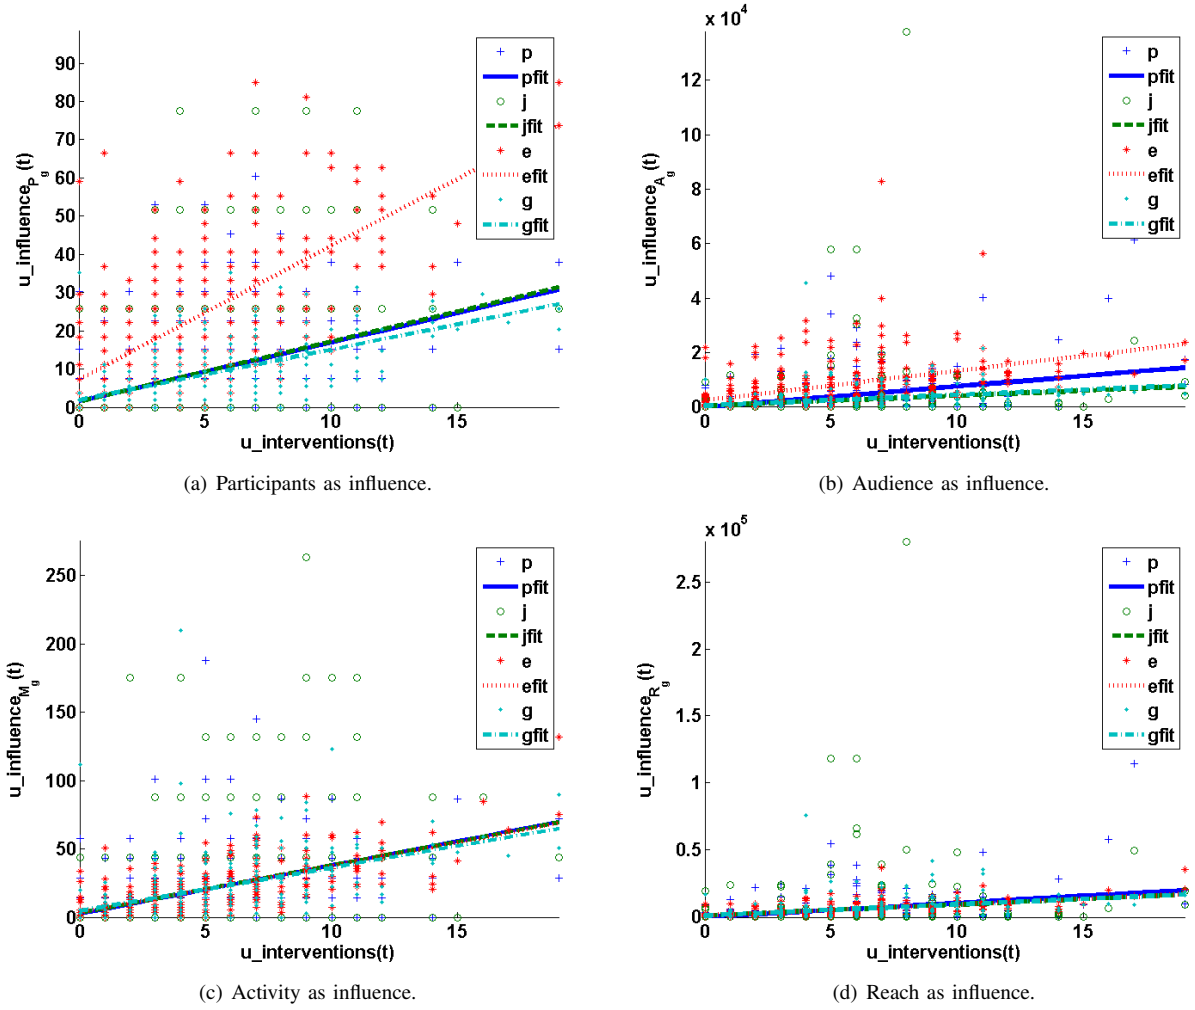

(a) Participants as influence.

(b) Audience as influence.

(c) Activity as influence.

(d) Reach as influence.

Fig. 1. Regression of each influence type on intervention scale for each target group.

TABLE II
POST-HOC REGRESSIONS FOR UNCONTROLLED INTERVENTIONS.

Hypothesis	Target group	Influence type	β	t	p
1	Politicians	Participants	.457	8.920	0^{**}
2	Journalists	Participants	.376	7.059	0^{**}
3	Employees	Participants	.718	17.919	0^{**}
4	General public	Participants	.662	15.356	0^{**}
5	Politicians	Audience	.410	7.850	0^{**}
6	Journalists	Audience	.149	2.618	$.009^{**}$
7	Employees	Audience	.482	9.569	0^{**}
8	General public	Audience	.381	7.151	0^{**}
9	Politicians	Activity	.441	8.535	0^{**}
10	Journalists	Activity	.345	6.381	0^{**}
11	Employees	Activity	.722	18.149	0^{**}
12	General public	Activity	.531	10.883	0^{**}
13	Politicians	Reach	.405	7.697	0^{**}
14	Journalists	Reach	.160	2.812	$.005^{**}$
15	Employees	Reach	.635	14.280	0^{**}
16	General public	Reach	.422	8.088	0^{**}

$^{**}p <= .05/4 = .0125$ (Bonferroni criterion per influence type).

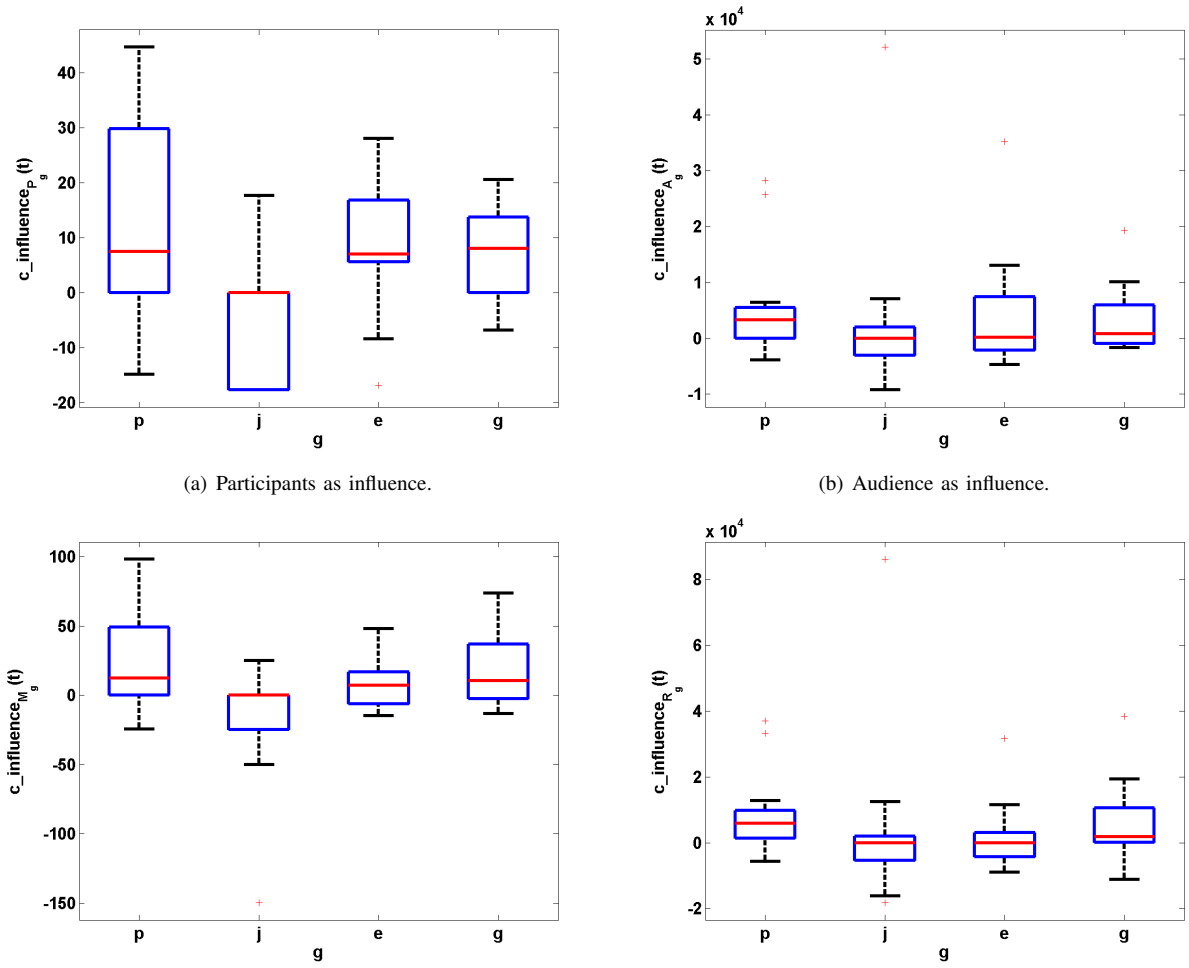

(a) Participants as influence.

(b) Audience as influence.

(c) Activity as influence.

(d) Reach as influence.

Fig. 2. Boxplots of the effect of controlled interventions on each influence type for each target group.

As can be seen in Table III, two-tailed independent t-tests (with $n = 14$ and $df = 13$) showed that the influence of controlled interventions was significantly higher than 0 for the general public, but only for the participants influence type. All other influences were not significant. Hyp. 1 could therefore not be accepted for controlled interventions for each influence type and target group, except for the general public combined with participants.

2) Comparing influences per target group: Four separate one-way analyses of variance (ANOVA) showed significant main effects of target group, for participants, $F(3, 52) = 3.62$, $p = .02$, and activity, $F(3, 52) = 5.11$, $p = .0036$, but not for audience, $F(3, 52) = .36$, $p = .78$, and reach, $F(3, 52) = .32$, $p = .81$. Tukey's honestly significant difference post hoc analyses were therefore only conducted for the participants and activity influence types.

a) Participants: For the participants, controlled interventions did not have more influence on employees than on the general public (no difference and stronger than on journalists). Influence on politicians was not weaker than on the general public (no difference and stronger than on journalists). Influence on journalists was not stronger than on the general public (no difference). Hyp. 2, 3 and 4 could therefore all not be accepted for controlled interventions and participants.

b) Activity: For the activity, controlled interventions did not have more influence on employees than on the general public (no difference). Influence on politicians was not weaker than on the general public (no difference and stronger than on journalists). Influence on journalists was not stronger than on the general public (it was weaker). Hyp. 2, 3 and 4 could therefore all not be accepted for controlled interventions and activity.

VI. Conclusions

This paper aimed to contribute to the field of information diffusion by introducing four measures of influence. In particular, we examined: 1) to what extent people are in general capable of making interventions on social networks influencing users in this network (uncontrolled interventions), and 2) to what extent large governmental organizations are capable of doing this (controlled interventions).

TABLE III
POST-HOC T-TESTS FOR CONTROLLED INTERVENTIONS.

Hypothesis	Target group	Influence type	μ	t	p
1	Politicians	Participants	11.183	2.407	0.032^{*}
2	Journalists	Participants	-4.688	-1.150	0.27
3	Employees	Participants	13.206	2.397	0.032^{*}
4	General public	Participants	4.994	2.917	0.012^{**}
5	Politicians	Audience	6712.3	2.010	0.066
6	Journalists	Audience	3592.5	0.784	0.45
7	Employees	Audience	6060	1.393	0.19
8	General public	Audience	1946.1	1.961	0.072
9	Politicians	Activity	6.970	2.727	0.017^{*}
10	Journalists	Activity	-24.375	-1.547	0.15
11	Employees	Activity	29.707	1.813	0.093
12	General public	Activity	20.729	2.751	0.017^{*}
13	Politicians	Reach	2084.4	2.713	0.018^{*}
14	Journalists	Reach	5616.9	0.640	0.53
15	Employees	Reach	9088.8	0.734	0.48
16	General public	Reach	4792.2	1.856	0.086

$^{*}p <= .05$.
$^{**}p <= .05/4 = .0125$ (Bonferroni criterion per influence type).

Results show a positive relation between uncontrolled interventions and all influence measures used for all target groups (general public, politicians, journalists and employees). Results also show that controlled interventions have less influence: Significant influence was found for the general public, only mildly significant influence was found for politicians and employees, and no influence was found for journalists (as compared to uncontrolled interventions where influence was found for all target groups).

In general, the above suggests that influence is very well reachable, for any of the targeted groups or used influence measure, provided the scale of interventions on Twitter is large enough. In our opinion, the fact that one actually *can* find an effect on influence for some of the targeted groups, in spite of the (by default) much smaller scale of controlled interventions as compared to uncontrolled interventions, is promising. To find out whether more extensive attempts of controlled interventions have a larger effect on influence still remains a question. This underlines the necessity to conduct similar research on controlled interventions from other organizations than Dutchorg as well, which would also provide more ground for gathering evidence for the stated hypotheses.

In addition to this, we found that interventions influence groups to a different extent. The uncontrolled interventions influenced all target groups, though these groups were affected to a different degree. Own employees were influenced strongest, differing significantly from the other groups.

In public opinion formation it makes sense that not everyone is equally receptive to interventions. Therefore people are influenced differently. The same logic applies to social networks. Clearly, organizations can influence social media users, especially in relatively small social networks. It is likely that substantial effects also depend on the number of people in the social network. This poses an interesting venue for future research (intervention scale vs. network size).

Finally, in the present study we did not analyze how the different influence types are related to each other. Doing so could provide more knowledge on the incremental validity of those measures. We furthermore did not look at the effect of interventions from different sources, other than the distinction "uncontrolled" and "controlled", such as those from the in this data already annotated employees, journalists, politicians and general public. Further research should focus on all of these matters as well.

REFERENCES

[1] B. A. Huberman, D. M. Romero, and F. Wu, "Social networks that matter: Twitter under the microscope," *First Monday*, vol. 14, no. 1, 2009.

[2] E. T. R. Rosenman, "Retweets — but not just retweets: Quantifying and predicting influence on twitter," 2012, bachelor's thesis, applied mathematics, Harvard College, Cambridge.

[3] J. Yang and S. Counts, "Predicting the speed, scale, and range of information diffusion in twitter," in *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*, 2010.

[4] L. Dugan. (2011) 5 tools to measure your twitter influence. [Online]. Available: http://www.mediabistro.com/alltwitter/5-tools-to-measure-your-twitter-influence_b3570

[5] R. Mehta, D. Mehta, D. Chheda, C. Shah, and P. M. Chawan, "Sentiment analysis and influence tracking using twitter," *International Journal of Advanced Research in Computer Science and Electronics Engineering*, vol. 1, no. 2, pp. 72–79, 2012.

[6] P. A. Grabowicz, J. J. Ramasco, E. Moro, J. M. Pujol, and V. M. Eguiluz, "Social features of online networks: The strength of intermediary ties in online social media," *PLoS ONE*, vol. 7, no. 1, 2012, e29358.

[7] 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*, 2012, pp. 519–528.

[8] K. Sassenberg and M. Boos, "Attitude change in computer-mediated communication: Effects of anonymity and category norms," *Group Processes & Intergroup Relations*, vol. 6, no. 4, pp. 405–422, 2003.
[9] C. Pornpitakpan, "The persuasiveness of source credibility: A critical review of five decades' evidence," *Journal of Applied Social Psychology*, vol. 34, no. 2, pp. 243–281, 2004.
[10] E. S. Knowles and J. A. Linn, *Resistance and persuasion*. Lawrence Erlbaum, 2003.
[11] R. Wijn and K. van den Bos, "On the social-communicative function of justice: The influence of communication goals and personal involvement on the use of justice assertions," *Personality and Social Psychology Bulletin*, vol. 36, no. 2, pp. 161–172, 2010.
[12] J. T. Cacioppo and R. E. Petty, "Persuasiveness of communications is affected by exposure frequency and message quality: A theoretical and empirical analysis of persisting attitude change," *Current issues and research in advertising*, vol. 3, no. 1, pp. 97–122, 1980.
[13] D. Sullivan. (2011) Why "second chance" tweets matter: After 3 hours, few care about socially shared links. [Online]. Available: http://searchengineland.com/why-second-chance-tweets-matter-after-3-hours-few-care-about-socially-shared-links-92125
[14] T. Meyer and L. P. Hinchman, *Media democracy: how the media colonize politics*. Polity Press Cambridge, 2002.
[15] C. Van der Eijk, "The netherlands: media and politics between segmented pluralism and market forces," in *Democracy and the Media: A comparative perspective*, R. Gunther and A. Mughan, Eds. Cambridge University Press, 2000, ch. 9, pp. 303–342.