On the (In)effectiveness of the Share/Tweet Button: Idea Management for Civic Participation

The authors study the practice of promoting idea campaigns in social networks via the well-known Share/Tweet button. They analyze data about 53 civic participation initiatives collected from IdeaScale, one of the leading online idea management platforms today, and unveil a considerable misconception about the effectiveness of the practice. The article highlights open challenges and suggests a set of alternative techniques to leverage on the ideation capacity of social networks.

Jorge Saldivar
Catholic University "Nuestra Señora de la Asunción," Paraguay

Carlos Rodríguez
University of New South Wales, Sydney, Australia

Florian Daniel
Polytechnic University of Milan, Italy

Fabio Casati
University of Trento, Italy and Tomsk Polytechnic University, Russia

Luca Cernuzzi
Catholic University "Nuestra Señora de la Asunción," Paraguay

Idea management (IM) is the process of collecting, developing, and selecting ideas to develop new, innovative products, services, or regulations, or to improve existing ones.[1] We call these processes IM initiatives. The practice isn't limited to commercial domains only and, boosted by the World Wide Web, has recently been gaining momentum in the domain of politics and civic participation open to the general public.[2] A prominent example of this practice took place in Finland in 2013 where the public participated in an off-road traffic law reform initiative. The Finns participated online in the lawmaking process by submitting their ideas and by commenting and voting on others' ideas.[3] Similar initiatives are emerging all over the world.[4,5]

The Finnish case emblematically shows how online IM in general has evolved from the naive "leave feedback" form of only a few years ago into dedicated applications and, more recently, full-fledged IM platforms. Examples of popular IM platforms are IdeaScale (http://ideascale.com), Crowdicity (http://crowdicity.com), and MindMixer (http://mindmixer.com).

To increase the visibility of initiatives and attract participants (members of the initiatives), increasingly these platforms leverage on social networks, such as Facebook and Twitter. This practice's effectiveness, however, isn't proven yet. In this article, we take a closer look

 Published by the IEEE Computer Society 1089-7801/17/$33.00 © 2017 IEEE

Research on Idea Management

Research on IM so far has focused on improving the methods used to define suggestions, the mechanisms used to display streams of ideas, the features applied to value proposals, the approaches employed to find contributions, and similar endeavors. Deliberation maps,[1] for instance, structure participants' contributions as problem trees containing the problem to solve, potential solutions, and arguments for and against proposed solutions.

The use of semantic technologies is proposed by Adam Westerski and colleagues[2] to organize, link, and classify the ideas using metadata annotations. Improving scoring methods used to value the ideas is the goal of Anbang Xu and colleagues,[3] who present a reference-based scoring model as an alternative to the traditional thumbs up/down voting systems.

Siamak Faridani and colleagues[4] introduce a two-dimensional visualization plane to address the filter bubble effect — narrowing the exposure to recent, popular, or controversial information — of linear lists used to display opinions in online sites. Gregorio Convertino and colleagues[5] address information overload in the evaluation phase with natural language processing methods to identify the core of proposals and to quickly discover reactions inside comments.

In a similar vein, Efthimios Bothos and colleagues[6] propose the application of information aggregation markets to facilitate the evaluation of ideas. Jennie Björk and Mats Magnusson explored the inclusion of social networks into ideation processes, analyzing the relationship between the quality of ideas and the connectivity (degree centrality) of the contributors.[7]

References

1. M. Klein and L. Iandoli. "Supporting Collaborative Deliberation Using a Large-Scale Argumentation System: The MIT Collaboratorium," *Proc. Directions and Implications of Advanced Computing*, 2008, pp. 5–12.
2. A. Westerski, C. Iglesias, and F. Tapia Rico. "A Model for Integration and Interlinking of Idea Management Systems," *Metadata and Semantic Research*, Springer, 2010, pp. 183–194.
3. A. Xu and B. Brian. "A Reference-Based Scoring Model for Increasing the Findability of Promising Ideas in Innovation Pipelines," *Proc. Conf. Computer-Supported Cooperative Work and Social Computing*, 2012, pp. 1183–1186.
4. S. Faridani et al., "Opinion Space: A Scalable Tool for Browsing Online Comments," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2010, pp. 1175–1184.
5. G. Convertino, A. Sándor, and M. Baez. "Idea Spotter and Comment Interpreter: Sensemaking tools for Idea Management Systems," *Proc. 6th Int'l Conf. Communities and Technologies: Workshop on Large-Scale Ideation and Deliberation Systems*, 2013; https://disi.unitn.it/,baez/files/sandor2013spotter.pdf.
6. E. Bothos, D. Apostolou, and G. Mentzas, "A Collaborative Information Aggregation System for Idea Management," *Proc. 3rd Int'l Conf. Internet and Web Applications and Services*, 2008, pp. 289–296.
7. J. Björk and M. Magnusson, "Where Do Good Innovation Ideas Come from? Exploring the Influence of Network Connectivity on Innovation Idea Quality," *J. Product Innovation Management*, vol. 26, no. 6, 2009, pp. 662–670.

at this approach; we're particularly interested in understanding the effectiveness of the common Share/Tweet button featured by most modern websites, including IM platforms, and articulate our research question into the following hypotheses:

- *H1*. Sharing/tweeting about civic participation initiatives in Facebook/Twitter increases the number of people registered as members of the initiatives.
- *H2–H4*. A higher sharing/tweeting activity per member leads to higher productivity in ideas (H2), votes (H3), and comments (H4) per member.

Here, we test these hypotheses by analyzing data about 53 publicly accessible civic participation initiatives from IdeaScale and report on our findings, also discussing open issues and alternative ways of accessing social network communities more effectively.

Dataset

Our dataset consists of public access innovation initiatives on IdeaScale, active as of March 2014. Organizers of IdeaScale initiatives define, as part of the setup process, a list of categories or campaigns inside which the community of participants can post their ideas. An idea is composed of a title and a description. Members of the community can comment and assign positive or negative valuations (votes) to others' ideas; they also can share the ideas within their social networks. Figure 1a introduces snapshots of one of the dataset's initiative's home page in IdeaScale and Figure 1b shows the voting, commenting, and social network-sharing features.

The dataset contains 73 idea management initiatives oriented to civic participation, of which 10 don't enable the Share and Tweet buttons – key elements for our study. Of the remaining 63, we excluded 10 more

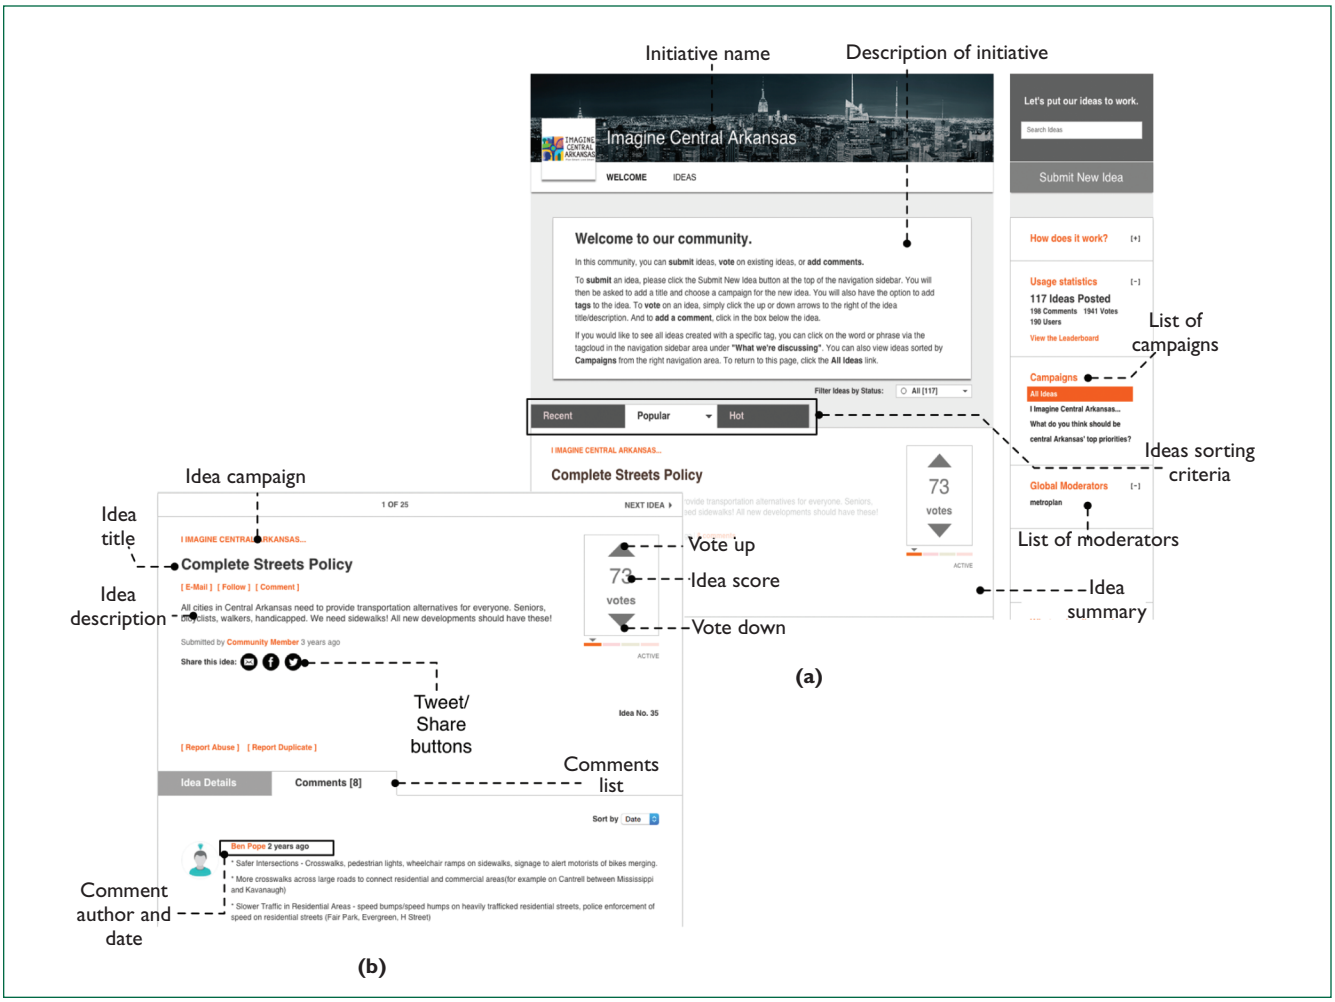

Figure 1. User interface of IdeaScale. (a) Snapshot of an initiative's homepage in IdeaScale. (b) Detailed view of an idea submission and the commenting, social network sharing, and voting functions.

(leaving 53), because of their outlier numbers of members, ideas, votes, comments, shares or tweets.

The vast majority of the initiatives, 42 out of the 53 (79 percent), engage citizens in discussions on topics of public interest. Almost half of the initiatives are sponsored by public institutions, such as the Helsinki Public Transportation Office, the US Patent and Trademark Office, or Redmond City Government. The goal is to harvest ideas from citizens on how public services and infrastructures (such as public transportation or downtown parks) or processes (such as a patent/trademark application process) can be improved. The rest of the initiatives are organized by civic organizations (Imagine Central Arkansas, CambiAnzio, and Public Works Agency), political associations (Manhattan Young Democrats, Politica Oltre, and Cinque Stelle Movement), or supported by ad hoc communities of citizens that gather together to exchange ideas on how their cities' services (garbage collection, connectivity, libraries, and parks) can be improved. The remaining 21 percent of the initiatives (11) are carried out by political and civic organizations that seek to involve their members in discussions about in-house topics. In the following, we refer to these two clusters as the *Public* and *In-house* clusters.

Together, all initiatives in our dataset account for 5,288 members and register 2,659 ideas — of which 55 are tagged as implemented or in progress of implementation — 22,332 votes and 3,855 comments. At the moment we collected the data, the initiatives and their ideas

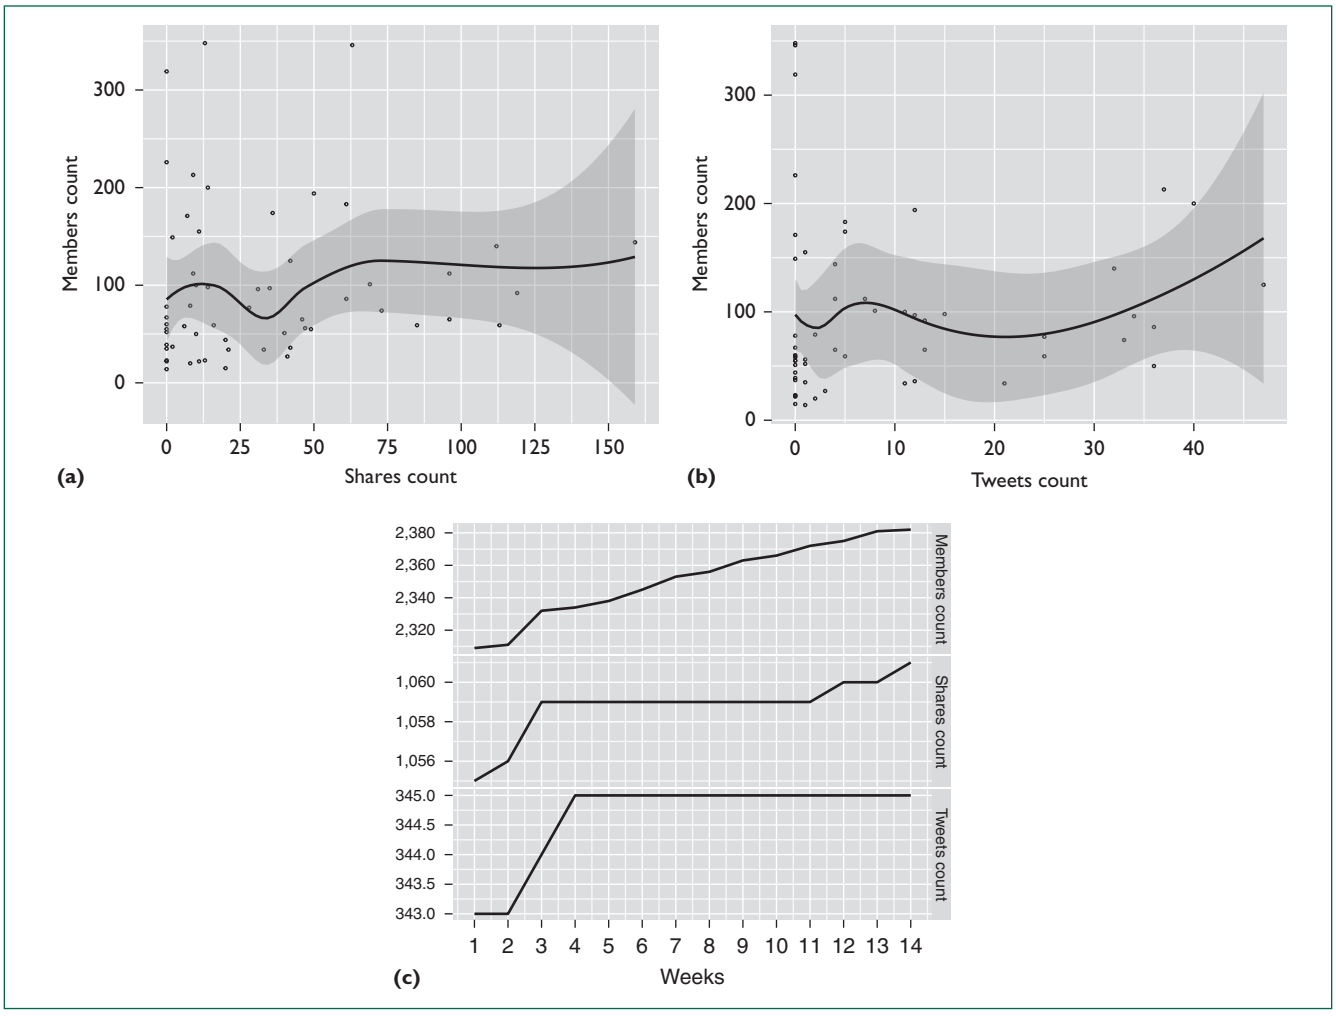

Figure 2. Correlation analysis of the number of shares/tweets and the number of members of 53 idea management (IM) initiatives, and evolution of shares, tweets, and members over 14 weeks for 26 active initiatives. (a) Shares versus members: a larger number of shares doesn't lead to a larger number of members. (b) Tweets versus members: tweeting doesn't lead to larger communities of members. (c) Longitudinal study: members grow faster than shares and tweets.

were promoted a total of 1,825 times on Facebook and 483 times on Twitter using the Share and Tweet buttons, respectively. Also, 49 percent (26) of the initiatives showed to be actively running, while 51 percent (27) didn't show activity in the last six months before March 2014. The biggest and most significant group of initiatives (the Public cluster) report 4,137 members and record 2,195 ideas — 54 marked as implemented or in process of the implementation — 18,426 votes and 3,519 comments, with 1,411 and 411 Facebook and Twitter shares, respectively. (The source code of the crawler, datasets, and R scripts of this study are available at http://github.com/joausaga/ims-sn-study.)

Enrollment of Members

We start our analysis to answer hypothesis H1 by scatterplotting the shares/tweets count versus the members count for the 53 initiatives (see Figures 2a and 2b). For an effective visualization, we also plot a Loess nonparametric regression curve[6] that fits the data points with a 95 percent confidence interval. It's immediately evident that the initiatives with higher sharing/tweeting activity aren't necessarily those with the larger numbers of members.

A pair-wise correlation analysis shows very low correlation (0.12 for members-shares and 0.05 for members-tweets), which unveils that, in general, increments in the number of shares/

tweets only unlikely affect positively the number of members. The situation doesn't change if we split the analysis by the identified clusters: 0.17 and −0.38 for members-shares and 0.17 and −0.34 for members-tweets in the Public and the In-house clusters, respectively. This, however, provides only a static picture of the data.

To obtain insight into the dynamics of the IM ecosystem and to understand whether shares/tweets help increase participants over time or whether increments are more due to the simple passing of time, we designed a longitudinal analysis for the 26 initiatives of the whole dataset that were effectively active at the time of our observation. Once a week from March to May 2014 (14 weeks) we recorded the number of members, shares, and tweets for these initiatives. Figure 2c depicts the identified evolution. The number of members grew over the 14 weeks of the study, passing from about 2,233 to more than 2,305 at the end of the study. Shares and tweets reported only slight increments of change, together with long periods of stability. The number of tweets increased by 2 (from 343 to 345) from weeks 2 to 4 and remained constant for the rest of the period. The number of shares grew from weeks 2 to 3 and stayed unaltered until week 11, when it increased again above 1,060 shares at the end of the study (starting from 1,055).

At this point, it appears to be clearer that increments in the number of shares/tweets are only marginally related with increments in the number of members. To quantify the real influence of shares/tweets and the initiatives' lifetime — that is, the elapsed time between the start of the longitudinal study and the end of it (in our case 14 weeks) on attracting members — we calculate for the 26 active initiatives the difference in members, shares, and tweets between the beginning and end of the observation period and conduct multiple regression analyses. Specifically, the relative impact of shares/tweets against lifetime is measured by two different regression analyses: one analysis including shares and lifetime as independent variables and another considering tweets and lifetime as the regression coefficients. In both cases, the variance of members (M) appears to be explained by the combination of these variables. Shares (S) and lifetime (L) account for 98 percent (F (2, 11) = 289.6, p-value < 0.05, $M = -0.001 + 3.17$ (p-value 0.01) $S + 0.035$ (p-value 9.25 $e - 8$) L) of the variance in members, while tweets (T) and lifetime have an impact of 99 percent (F (2, 11) = 4,269, p-value < 0.05, $M = -0.030 + 6.59$ (p-value 0.001) $T +$ 0.036 (p-value 2.26 $e - 9$) L). A comparison of the relative importance of the variables unveils that the lifetime explains the largest amount of the variance compared to Facebook shares and tweets (62 percent and 67 percent, respectively). A finer-grained regression analysis limited only to the initiatives that showed social activity during the period of observation (all part of the Public cluster) reports a similar trend: that about 65 percent of the member variation is explained by the initiatives' lifetime.

The evidence collected via both the correlation analysis and the regression analysis doesn't provide enough arguments to accept hypothesis H1 that sharing/tweeting increases the number of members of idea management initiatives.

Ideation Productivity

Next, we study whether the social networking activity of members impacts the amount of ideas, votes, and comments produced by the initiatives.

A factor that might affect the production of ideas, votes, and comments is, of course, the number of members of the initiatives: you might guess that the more participants an initiative has, the more ideas, comments, and votes you can expect. Suitable correlation analyses on these variables confirm that the number of members is indeed significantly and positively correlated with the number of ideas ($r = 0.64$, $p < 0.05$), votes ($r = 0.67$, $p < 0.05$), and comments ($r = 0.43$, $p < 0.05$).

To diminish the bias introduced by the number of members in the study of the impact of sharing/tweeting, we proceed our analysis with the relative numbers of ideas, votes, and comments per member (productivity per member). That is, we measure whether the ratios of shares/tweets over members influences the productivity of ideas, votes, and comments of the initiatives and study hypotheses H2–H4.

The scatterplots in Figure 3 reveal that many Facebook shares or Twitter tweets per member don't necessarily lead to higher productivity. Interestingly, the most productive initiatives seem to have scarce tweeting activity per member, while for the Facebook shares per member,

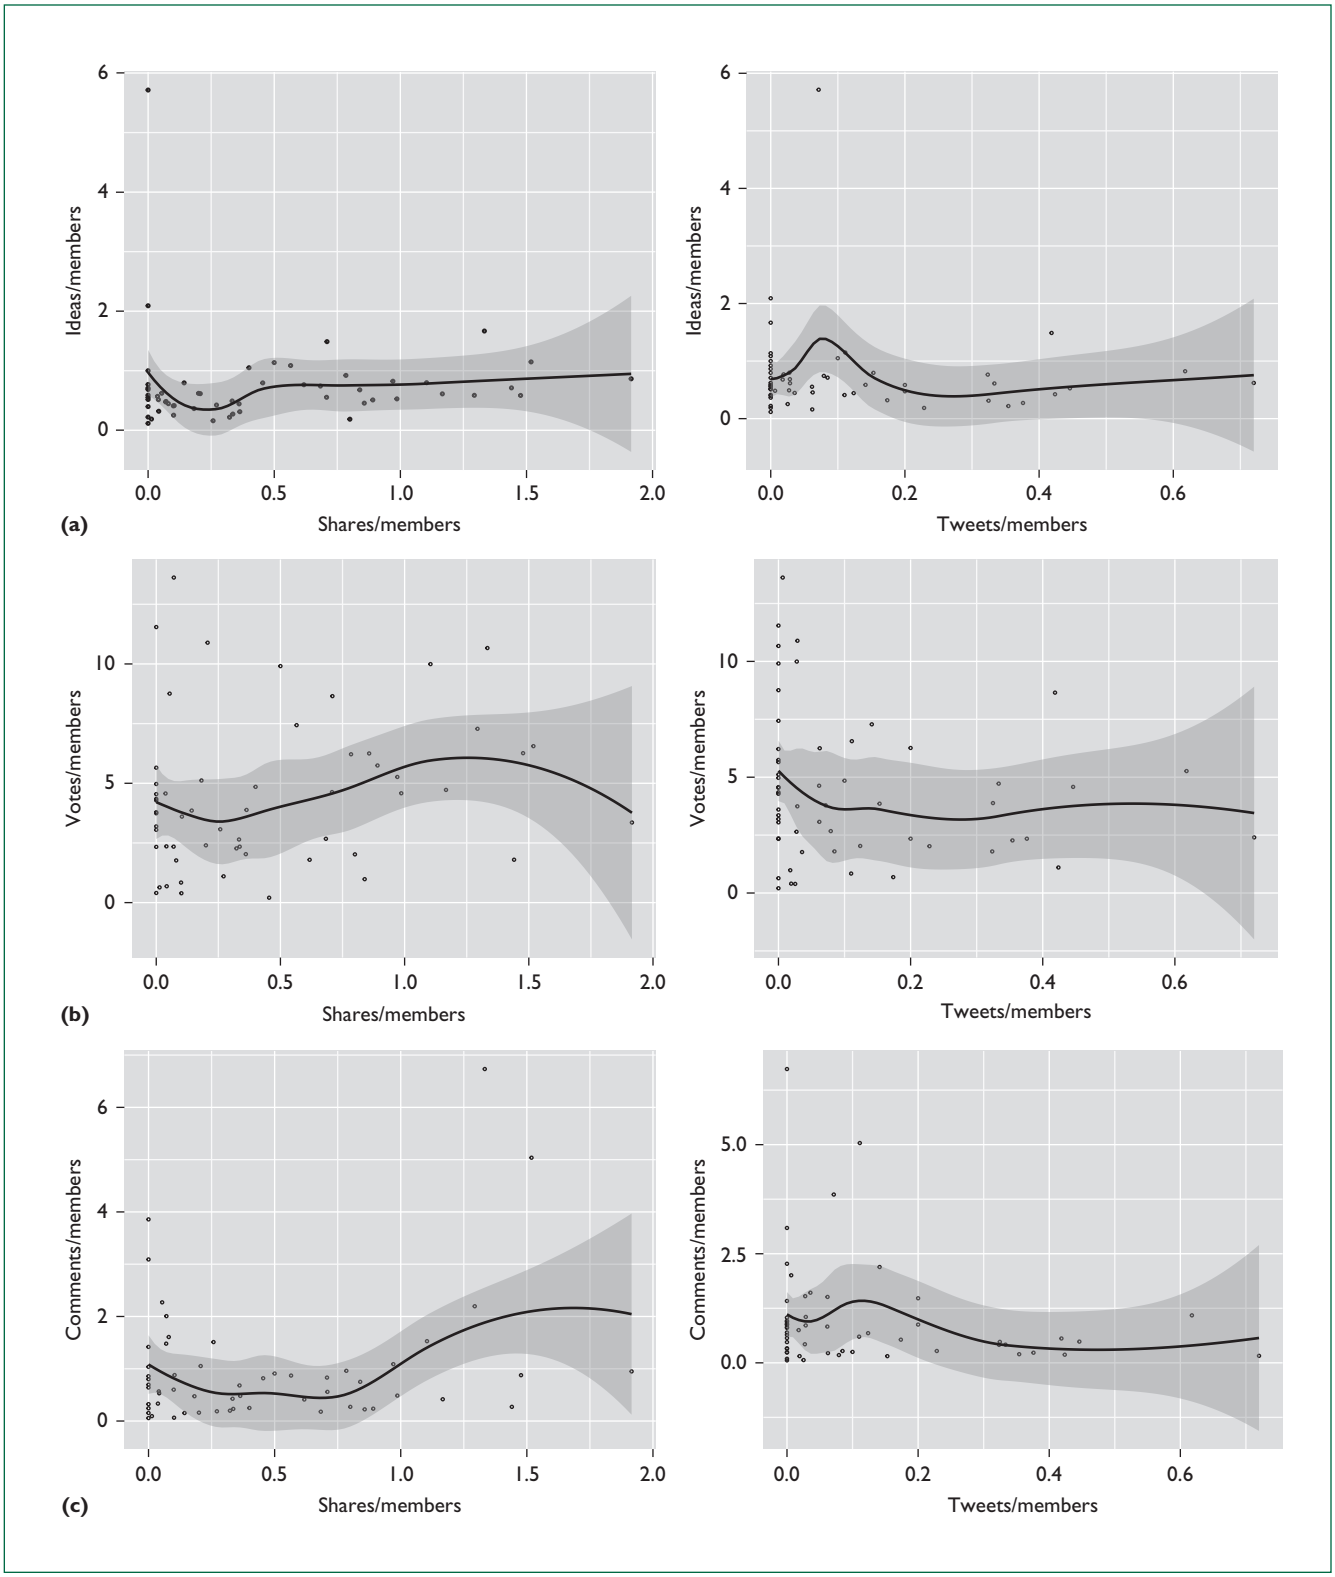

Figure 3. Correlation of the productivity of members (ideas, votes, and comments per member) and the average member's social networking activity (shares/tweets per member). (a) Shares/tweets per member versus ideas per member. (b) Shares/tweets per member versus votes per member. (c) Shares/tweets per member versus comments per member.

each plot has its own dynamic. Figure 3a shows that most initiatives have only small values of shares/tweets per member, highlighting that the productivity of ideas is almost unrelated to sharing/tweeting. As for the votes, Figure 3b (left) shows a slight increase in the productivity for share ratios between 0.5 and 1.5. It appears that the number of shares per members affects the productivity of votes when at least one share is generated for every two members. As for the comments, the left plot of Figure 3c seems to indicate that the ratio of shares over members positively contributes to the productivity of comments as soon as the members produce at least one share on average.

We also analyze the correlation on these variables. The number of shares per member is only slightly correlated with the number of ideas ($r = 0.03$, p-value $= 0.84$), votes ($r = 0.20$, p-value $= 0.15$), and comments per member ($r = 0.24$, p-value $= 0.08$). The number of tweets per member also has a low dependence on the number of ideas ($r = -0.05$, p-value $= 0.74$), votes ($r = -0.13$, p-value $= 0.35$), and comments ($r = -0.18$, p-value $= 0.21$) per member. These numbers confirm analytically what was anticipated intuitively by the plots in Figure 3: the productivity of ideas, votes and comments seems to be independent of the sharing and tweeting activity of the initiatives' members.

Similar low correlations also hold for the Public and In-house clusters individually. An interesting exception can be identified for the In-house cluster, where sharing on Facebook has a positive influence on the number of ideas per member ($r = 0.68$, p-value $= 0.02$). This correlation is likely explained by the tighter relationship that binds the members of an organization: they know each other, and many of them are also friends on Facebook. This is fundamentally different from the general audience targeted by the Public cluster.

In summary, we thus accept hypothesis H2 for the In-house cluster limited to Facebook shares and idea productivity, while we reject hypotheses H2–H4 for the Public cluster in general and the other combinations studied for the In-house cluster.

Ideation Inside Social Networks

Given the aforementioned results, next we try to understand in more detail what happens when information about IM initiatives is promoted inside social networks using the Tweet button and whether social networks are suitable at all for IM. We limit our analysis to Twitter, as the majority of its content is publicly accessible (99 percent according to Mashable's social media expert Kurt Wagner: http://mashable.com/2013/08/13/topsy-opens-twitter-data). This differs from Facebook, because Facebook's posts are strongly regulated by privacy policies and generally not publicly accessible.

Usually, the Tweet button is equipped with a default message that prefills the Compose box of tweets. Because the goal of tweeting is to drive traffic to an initiative's website, this default message typically contains the website's URL, among other properties. We can use this URL as an identifier: using the REST API of Twitter and Topsy (a social analytics service that existed at the moment this article was written but now is defunct), we searched for the URLs of the initiatives' websites as well as for the URLs of their ideas (in IdeaScale, every idea is accessible through a dedicated URL). We collected a total of 723 tweets, of which 265 are about initiatives and ideas posted via the Tweet button, whereas the remaining 458 tweets were posted using other means, such as Twitter's WebClient, smartphone app, or other external clients, such as Buffer, TweetAdder, or Hootsuite. The vast majority of tweets (81 percent) was published by members; if we match the tweets' handlers with the username of moderators and administrators or with the name of the initiatives, we see that the remaining 19 percent of the tweets were authored by the initiatives' organizers.

A manual inspection of a sample of the collected tweets unveiled that members use Twitter for generating awareness (in line with its use in general), as the following example shows: "We want to hear your ideas! #transformrockford" (@TransformRkfd) and "Do you have an idea for Huntsville? Join the discussion at Imagine Huntsville http://www.imaginehuntsville.com" (@HSVevents). However, here Twitter serves two specific purposes: to promote ideas and fuel the discussion; and to cast votes for ideas. An instance of these purposes can be found in the following tweet that promotes an idea and requests voting actions from followers: "This is awesome, guys. Pls RT & Vote for the game Myopia in the @WhiteHouse Initiative Games For Impact http://gamesforimpact.ideascale.com

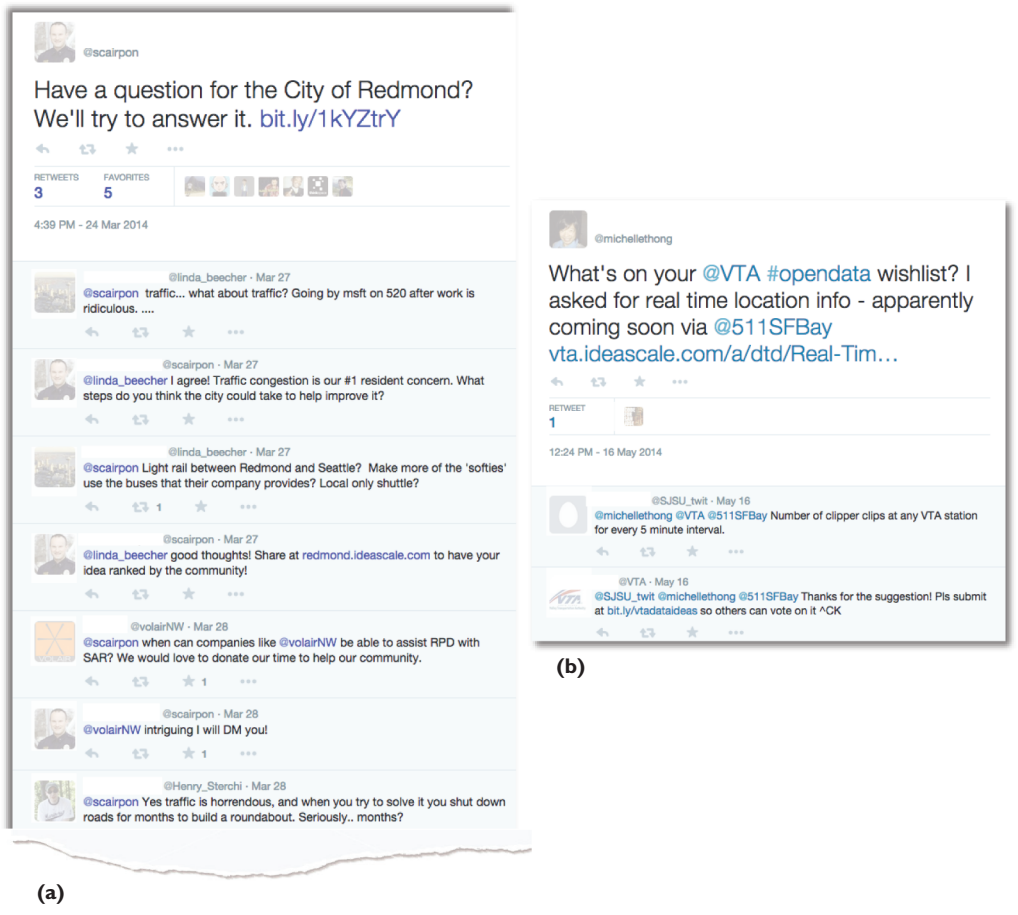

Figure 4. Two examples of manually written tweets with an excerpt of the value-adding reactions they triggered. (a) Discussion about initiative "City of Redmond." (b) Suggestion posted by @SJSU_Twitt within the context of the initiative "VTA."

/a/dtd/MYOPIA-An-intergenerational-collective-action-game-series" (@jesserker). Through this analysis, it was discovered that moderators' tweets target similar goals: create awareness, promote interesting ideas, cast votes for ideas, and publicly thank members for their contributions.

The Tweet button's effectiveness can be gauged by comparing the reactions its tweets raised against the reactions triggered by the tweets coming from others sources (reactions are measured by summing up the number of retweets, replies, and favorites). The data we collected show that tweets generated with the Tweet button produced on average about three times fewer reactions: tweets posted using the Tweet button triggered on average 0.39 reactions, while tweets published through other means raised on average 1.30 reactions. Moreover, with a 95 percent of confidence ($p < 0.05$) we can say that the average number of reactions triggered via Twitter's WebClient and other clients is higher by two to three times (0.70 to 1.12). The maximum number of reactions triggered by tweets posted through the Tweet button is seven, whereas tweets published using other Twitter clients received from three to even about 30 reactions. Repeating the same analysis for moderators/administrators and members individually doesn't reveal any difference among the two types of participants.

Our intuition is that the difficulty to gain attention with the Tweet button might be the fruit of its generic and impersonal nature (default text only). In contrast, tweets posted through other means are usually written manually and contain personal comments, emotions, excitement, or something similar — these are

all characteristics that automatically generated tweets don't have.

For instance, in Figure 4 we present a couple of interesting tweets worth noting. Figure 4a introduces a sample of messages exchanged between the followers of @scarpon (the moderator of the initiative City of Redmond) about improving the public services of Redmond, Washington. The long discussion produced 36 tweets from 20 different participants and generated valuable content, which probably wasn't transported back to IdeaScale and, hence, lost. In fact, Figure 4b captures a case where a Twitter user contributed to the initiative called "VTA," triggering the answer "Thanks for the suggestion! Pls submit at http://vta.ideascale.com so others can vote on it" (@VTA). The suggestion was considered just as valuable as suggestions generated within the "official" platform. However, unless the moderator moves the content of the tweet to IdeaScale or the person who posted it takes the time to do so, the contribution runs the risk of getting lost. Losing this kind of feedback could be a huge loss. It suffices to recall that Iceland's citizens employed Facebook, Twitter, YouTube, and Flickr to reform their national constitution.[4]

The findings we report on in this article somewhat surprisingly reveal that the Share/Tweet buttons are, in general, not effective at helping IM platforms to increase participation or productivity. However, they might work in situations where the members are already connected through online social relationships, such as the case of the initiatives in the In-house cluster. It's evident that social networks have a huge potential as incubators of ideas and proposals, yet current techniques fail to leverage this properly. In fact, even if triggered by Facebook shares or Twitter tweets, people inside social networks apparently aren't willing to go to and register for another platform, and don't allow IM initiatives to track and value their ideas and feedback.

We're aware that these findings are specific to the context of idea management for civic participation and limited by the study's observational nature (for example, we couldn't test reactions to artificial stimuli). Also, the study might suffer from "lurking" variables, such as unattractive discussion topics, uncommitted organizers or moderators, unclear participation rules, or the timing of our observation (we couldn't study the startup phase of new initiatives). However, the study provides an analytical picture of a domain that has strong commonalities with other contexts that aim to attract people from social networks to their own platform, application, or initiative (such as advertisement or entertainment).

The challenge seems to be how to harvest the ideas and feedback people leave inside social networks. This is an engineering problem that, first and foremost, requires understanding and leveraging existing social network usage conventions. In the specific context of IM, we identify three levels of intrusiveness of possible engineering approaches:

- *Using existing conventions*. This approach aims to identify ideation initiatives inside social networks, such as conversations among people, and to harvest ideas and feedback without touching the social networks themselves. An example is sentiment analysis.[7]
- *Introducing new conventions*. This approach aims to establish ideation-specific conventions, such as dedicated hashtags and conversation rules, to trigger ideation initiatives and to facilitate harvesting results. An example is the initiative MyIdea4CA, which was launched by former California governor Arnold Schwarzenegger to encourage citizens to post ideas for the state on Twitter with the hashtag #myidea4ca.[8]
- *Change of conventions*. This approach aims to introduce new features and conventions into social networks — for example, via functional extensions. An example is supporting the crowdsourcing of tasks inside social networks.[9]

Which of these approaches (or which combination of these approaches) performs best is something that still needs to be studied. However, as hinted by the findings of our study, a departure from the naive Share/Tweet buttons is a promising step forward that goes far beyond the domain of IM for civic participation.

Acknowledgments

This work was supported by the European Commission (project BPM4People) and by the project "Evaluation and

enhancement of social, economic and emotional wellbeing of older adults" under agreement 14.Z50.310029, Tomsk Polytechnic University.

References

1. M. Flynn et al., "Idea Management for Organisational Innovation," *Int'l J. Innovation Management*, vol. 7, no. 4, 2003, pp. 417–442.
2. S.-J. Min, "Online vs. Face-to-Face Deliberation: Effects on Civic Engagement," *J. Computer-Mediated Comm.*, vol. 12, no. 4, 2007, pp. 1369–1387.
3. T. Aitamurto and H. Landemore. "Democratic Participation and Deliberation in Crowdsourced Legislative Processes: The Case of the Law on Off-Road Traffic in Finland," *Proc. 6th Conf. Communities and Technologies Workshop: Large-Scale Idea Management and Deliberation Systems*, 2013; http://cochette.xrce.xerox.com/comtech13/papers/paper1_aitamurto_landemore.pdf.
4. H. Landemore. "Inclusive Constitution-Making: The Icelandic Experiment," *J. Political Philosophy*, vol. 23, no. 2, 2015, pp. 166–191.
5. S. Nambisan and P. Nambisan. *Engaging Citizens in Co-Creation in Public Services: Lessons Learned and Best Practices*, tech. report, IBM Center for the Business of Government, 2013.
6. W. Cleveland. "Robust Locally Weighted Regression and Smoothing Scatterplots," *J. Am. Statistical Assoc.*, vol. 74, no. 368, 1979, pp. 829–836.
7. B. Liu, *Sentiment Analysis and Opinion Mining*, Morgan & Claypool, 2012.
8. F. Comunello. *Networked Sociability and Individualism: Technology for Personal and Professional Relationships*, IGI Global, 2011, p. 241.
9. A. Bozzon, M. Brambilla, and S. Ceri. "Answering Search Queries with CrowdSearcher," *Proc. World Wide Web Conf.*, 2012, pp. 1009–1018.

Jorge Saldivar is a postdoctoral researcher at the Catholic University "Nuestra Señora de la Asunción," Paraguay. His research interests include crowdsourcing, collective intelligence, idea management, and citizen engagement. Saldivar has a PhD in information technology from the University of Trento, Italy. Contact him at jorge.saldivar@uc.edu.py.

Carlos Rodríguez is a postdoctoral researcher at the University of New South Wales, Sydney, Australia. His research interests include crowdsourcing, business process management, and big data analytics. Rodríguez has a PhD in information and communication technologies from the University of Trento, Italy. Contact him at carlos.rodriguez@unsw.edu.au.

Florian Daniel is an assistant professor at the Polytechnic University of Milan, Italy. His research focuses on crowdsourcing, web engineering, mashups, service-oriented computing, and business process management. Daniel has a PhD in information technology from the Polytechnic University of Milan. Contact him at florian.daniel@polimi.it.

Fabio Casati is a professor of social informatics at the University of Trento and at the Tomsk Polytechnic University, Russia. His work focuses on a research line of technologies for happiness, delivering results that have a direct impact on people's life. The research results are available at lifeparticipation.org. Casati has a PhD in information technology from the Polytechnic University of Milan, Italy. Contact him at fabio.casati@unitn.it.

Luca Cernuzzi is a software engineering professor at the Catholic University "Nuestra Señora de la Asunción," Paraguay. His current research interests include web engineering, data science, and information and communication technologies for social good. Cernuzzi has a PhD in engineering from the University of Modena e Reggio Emilia. Contact him at lcernuzz@uc.edu.py.