

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/301547682>

Using Online Controlled Experiments to Examine Authority Effects on User Behavior in Email Campaigns

Conference Paper · July 2016

DOI: 10.1145/2914586.2914619

CITATIONS

5

READS

172

4 authors, including:



Kwan Hui Lim

University of Melbourne

44 PUBLICATIONS 428 CITATIONS

[SEE PROFILE](#)



Ee-Peng Lim

Singapore Management University

486 PUBLICATIONS 9,936 CITATIONS

[SEE PROFILE](#)



Palakorn Achananuparp

Singapore Management University

61 PUBLICATIONS 703 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Bioinformatics [View project](#)



Social Network Analysis & Mining [View project](#)

Using Online Controlled Experiments to Examine Authority Effects on User Behavior in Email Campaigns

Kwan Hui Lim^{*†}, Ee-Peng Lim[†], Binyan Jiang^{#‡}, and Palakorn Achananuparp[†]

^{*}The University of Melbourne, Australia, and [†]Singapore Management University, Singapore

[#]Hong Kong Polytechnic University, Hong Kong, and [‡]Carnegie Mellon University, USA

limk2@student.unimelb.edu.au, {eplim,palakorna}@smu.edu.sg, binyanj@andrew.cmu.edu

ABSTRACT

Authority users often play important roles in a social system. They are expected to write good reviews at product review sites; provide high quality answers in question answering systems; and share interesting content in social networks. In the context of marketing and advertising, knowing how users react to emails and messages from authority senders is important, given the prevalence of email in our everyday life. Using a real-life academic event, we designed and conducted an online controlled experiment to determine how email senders of different types of authority (department head, event organizer and a general email account) affect the range of response behavior of recipients, which includes opening the email, browsing the event website, and registering for the event. In addition, we proposed a systematic approach to analyze the user response behavior to email campaigns from the time the user receives the email till he/she browses the website in a seamless manner.

Keywords: Online Controlled Experiments, A/B Testing, Behavioral Analysis, Authority, Email, Marketing

1. INTRODUCTION

Email is an important communication media, used by more than 2.58 billion users worldwide, who collectively generated more than \$13.6 billion in revenue [24]. In particular, individuals and organizations frequently use email for personal communication and internal coordination, respectively [19]. Many organizations use email not only for internal communications, but also in their marketing campaigns due to the prevalence of email and its low cost. From the perspective of an organization, email campaigns provide a good return on investment as “for every \$1 spent, \$44.25 is the average return on email marketing investment” [25].

Fig. 1 outlines the typical workflow of an event promotion campaign, which starts with the broadcasting of emails and ends with the user performing a specific action, e.g., register for an event or buy a product. There are however multiple

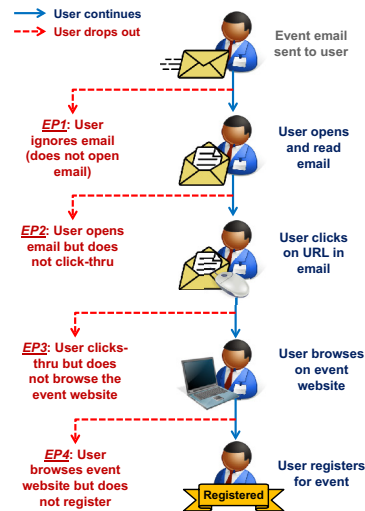


Figure 1: Typical Email Campaign Flowchart

exit points (EP) in the series of user actions, i.e., EP1-4 in Fig. 1. The user may possibly give up before achieving the desired end goal of registering for the event. To maximize the effectiveness of such campaigns, it is of utmost importance to understand the churn of users at each of these exit points so as to improve the design of future campaigns.

There is a large body of literature on the different factors that affect how likely a user will respond to an email [6, 17, 16, 23, 10]. Research has shown that authority is a key influence in people’s decision making process [14, 8]. In this research, we investigate the effects of two different authority status of senders: *superior* and *domain expert*, on user behaviors in email campaigns. The first type of authority refers to someone who is a *superior* in the organization hierarchy. The second type of authority refers to someone who is a *domain expert* with knowledge and experience not easily found among others. Conducting online experiments on users in a real world setting (as opposed to a laboratory environment) adds further complexity to this research, in the form of business processes imposed by the event organizer. The event and users are real and any experiments should not undermine the success of the event itself. Hence, cross-platform individual level behavior analysis, and organic experiment setting are the unique considerations in our work.

Research Objectives. In this online controlled experiment, we are interested in investigating the outcome of email campaigns and how it could be influenced by the authority status of email senders. An email campaign typically advertises a particular event, product or service with the goal of

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.

HT '16, July 10-13, 2016, Halifax, NS, Canada

© 2016 ACM. ISBN 978-1-4503-4247-6/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2914586.2914619>

motivating the user to comply with the marketing messages, e.g., register for the particular event or purchase/adopt the product/service. We measure the outcomes of cross platform email campaigns in terms of: (i) the open rate and click-through rate of emails; (ii) user activities on the event website; and (iii) registration rate of the event.

Contributions. Our main contributions are as follows:

- We propose a systematic approach to analyze email campaigns, with a focus on better understanding the potential exit points of users. (Fig. 1 and Sect. 4)
- We design a series of experiments to determine the effects of email senders’ authority status on their recipients’ behaviors in opening the email, browsing the event website, and registering for the event. (Sect. 3)
- We derive our key findings from the experiments, and show that the emails from superiors are highly effective in prompting users to react to the email and subsequently register for the event. (Sect. 4)
- We develop a multi-platform online experimentation system that facilitates the experimenters to seamlessly conduct an online controlled experiment over email and website. (Fig. 2 and Sect. 3.2)

2. RELATED WORK

There exists a large body of email-related studies that cover email response behavior [28], email mailbox management [29, 11, 27, 13, 12], email activity prediction [9, 22, 19, 5], and social network based email functions [30, 18]. Among these works, we are most interested in those related to email campaigns, in particular campaigns that involve elements of a controlled experiment. For example, [15] studied authority senders in the form of university professors and college students. Their experiment involved sending email requests (using the professor’s email and college student’s email) to a group of university students and random email addresses to complete an online survey. Similarly, [4] used the email accounts of a university’s Department of Transportation and Survey Center to send out email invitations for an online survey related to transportation issues. In contrast, [26] studied how authority logos and trust messages affect a user’s attitude towards online purchases using credit cards. They sent out messages with different authority logos (VeriSign, an international brand, and TWCA, a Taiwan-based domestic brand). Similarly, [7] studied the influence of sender and advertiser trust on user response in viral email advertising. They sent out viral marketing emails using trusted and less trusted email senders, where the emails contain content regarding trusted and less trusted advertisers.

As far as we are aware, there has been no prior work that concurrently study the authority effect of both the *superior* and *domain expert* on the entire spectrum of user response data (from email open to event registration), within the context of a real-life event (an academic workshop). While these earlier works highlight various interesting findings, our proposed research differs in the following ways: (i) instead of using an hypothetical scenario, our controlled experiment is conducted alongside a real-life event (i.e., an academic workshop), with the additional challenge of an experimental design that does not undermine the success of the event itself; (ii) instead of examining only a specific subset of aggregated data (e.g., only email click-through rate and time), we examine the entire spectrum of user response data (i.e., email




Figure 2: Multi-platform Online Experimentation System in Support of Experiment Design

open rate/time, email click-through rate/time, website visit, event registration, etc), thus providing a fine-grained analysis into user behavior from start to end; and (iii) instead of examining only a single type of authority, we concurrently examine the effects of both types of authority: *superior* and *domain expert*, within the same real-life event.

3. EXPERIMENTAL DESIGN

Our experiment was based on a two-day academic workshop that was attended by more than 150 participants, mainly from universities, research institutes and industry. This workshop was co-organized by three universities and included presentations by multiple speakers in the general area of data mining. Our experiment participants comprise 730 users, who belong to various departments from three large universities. The occupations of these users range from undergraduate and postgraduate students to university staff, representing the various occupations and hierarchical levels within a typical university.

3.1 Overview of Experiment

For this experiment, we define three types of email senders, namely: (i) general (non-personal) account; (ii) department superior; and (iii) event organizer. The general account is the email account of the organization hosting the event, while the latter two are the personal email accounts of Head of Departments (HOD) and event organizer, who represent the *superior* and *domain expert* authority respectively. Correspondingly, we want to identify three groups of users:

- Control Group (CGroup): Users who receive the event emails from a general account.
- Treatment Group 1 (TGroup1): Users who receive the event emails from their HODs.
- Treatment Group 2 (TGroup2): Users who receive the event emails from the event organizer.

For TGroup1, the HOD is the *superior* of the department that TGroup1 users currently belong to. For TGroup2, the event organizer acts as the *domain expert* in the area of data mining (the workshop topic) but otherwise have a minimal working relationship with TGroup2 users. CGroup reflects the mode of operation currently used by most organizations, i.e., they use a general organization email account to advertise their organization’s promotions, events and newsletters. In contrast, TGroup1 and TGroup2 represents the different type of authority figures with TGroup1 receiving emails from a *superior* (department head) while TGroup2 receives emails from a *domain expert* (event organizer). Table 1 shows the breakdown of users in their respective experiment groups.¹

¹Due to technical problems, we were unable to properly track the emails sent to some users in TGroup1. These users were excluded to maintain the accuracy of our experiments.

Table 1: User Breakdown of Experimental Groups

Group	Uni. A		Uni. B			Uni. C		Total
	Dept. A1	B1	B2	B3	C1	C2		
TGroup1	97	33	-	23	-	-	153	
TGroup2	98	35	76	24	30	25	288	
CGroup	97	35	78	24	30	25	289	
<i>Total</i>	<i>292</i>	<i>103</i>	<i>154</i>	<i>60</i>	<i>50</i>	<i>72</i>	<i>730</i>	

All users are uniquely identified by their email addresses and our experiment involves first sending emails to these 730 users, followed by tracking and collecting their activities on both the email (opens and click-throughs) and website (clicks and mouse-overs on page elements) platforms. In particular, we study the difference in user activities between users in CGroup (*general account*) and those in TGroup1 (*department head*) or TGroup2 (*event organizer*), by examining whether there are differences in terms of: (i) email open rate (Sect. 4.1); (ii) distribution of email open times (Sect. 4.2); (iii) email click-through rate (Sect. 4.3); (iv) distribution of email click-through times (Sect. 4.4); (v) probability of active users on the event website (Sect. 4.5); and (vi) probability of user registration (Sect. 4.6).

3.2 Multi-platform Online Experimentation System

To study user behavior across the different stages of an email campaign, we need to design a multi-platform online experimentation system that can automate and track such an experiment across both email and website platforms. Fig. 2 illustrates our proposed experimentation system, which comprises the following features:

1. Configure various parameters (e.g., user groups, email accounts, website URL) of the online experiment.
2. Allocate users to control and treatment groups based on random assignment.
3. Send emails to users in each group using a particular type of sender account, i.e., department head, event organizer, or general account.
4. Track user activities on the emails, namely email opens and user clicks on a link in the email, which then directs users to the event page.
5. If the user clicks-through to the event website, continue to track his/her activities on the website, i.e., clicks and mouse-overs on the various page elements.
6. Analyze the collected data of user activities on both email and website platforms.

While there exist applications for monitoring email campaigns and website visits [1, 2, 3], these applications monitor emails and websites independently, instead of as an integrated application. Even with integration, they only report activities at an aggregated level and do not provide analysis at the individual user level. These fine-grained user activities allow us to better understand why users leave at each exit point. More importantly, we would require a system that is able to conduct an integrated controlled experiment on both email and website and thereafter track the user activities as they traverse across these two platforms. Thus, we designed and developed a multi-platform online experimentation system that integrates with and extends our existing website experimentation platform, the **Living AnalyticS ExpeRimentations** system [21, 20]. This system enables us to easily conduct experiments to study user behavior across

the different stages of an email campaign. In addition, our proposed system automates and manages the main steps of this experiment, which we describe next.

3.3 Formatting and Sending of Emails

Our experimentation system facilitates the sending of emails to users in the three experimental groups. Before an email is sent, the system makes two main modifications to the email, namely: (i) it converts all URLs into a unique link that enables us to track when it is clicked; and (ii) it inserts into the email a tracking pixel that allows us to determine when the email is opened. In addition, we are able to track the user who has opened an email or clicked on an email link, and the time when this action was performed. From an experimenter’s point-of-view, this is akin to the broadcasting of an email and requires no additional workload on the experimenter apart from defining the experimental user groups and type of emails to send to each group.

We control for email content and design by using the same template for sending to the three experimental groups. At the top of this email template, the following text is added “FYI, you may be interested in this workshop. Best regards, [Sender Name]”. In addition, all emails are sent at approximately the same time to control for the effects of day and time on the three experimental groups. The main difference is the type of sender account and signing-off name that is used, i.e., HOD, event organizer or general account. This experiment design allows us to best evaluate the authority effects of email senders without the influence of the email content design and email delivery time.

3.4 Tracking of User Activities

As mentioned in the previous section, our system modifies each email before sending it out, and each reformatted link is unique to each user for a particular experiment. Once this link is clicked, it makes a web service call to our system that records the link which was clicked, the user who clicked it and the time of link clicking. Similarly, when the email is opened, the embedded tracking pixel makes a similar web service call to our system that records the user who opened this email and the time of email opening. To protect the privacy of users, the replies of users to the senders are directed to the true senders instead of our experimentation system.

One key functionality of our proposed system is its ability to automatically track users’ activities across multiple platforms, without the need for any intervention by the experimenter. If a user clicks on the email link to visit the event website, our system establishes a “handshaking” process, which establishes the email recipient as the website visitor. This “handshaking” is facilitated by the unique link created by our system during the email sending stage, which then enables us to uniquely identify this user by his/her email address as the user browses the website. Thereafter, we will be able to track and monitor the user’s activities (mouse clicks and mouse-overs) on the website. The tracking of user’s activities on both the email and website is automatically performed by our experimentation system, without needing the experimenter to develop any specific tracking program for the emails and websites.

4. USER BEHAVIOR ON EMAIL / WEBSITE

As illustrated in Fig. 1, there are four main exit points in a typical campaign where a user might possibly leave without

Table 2: Summary of Emails Sent, Opened and Clicked-through, and Active and Registered Users.

	Emails Sent	Emails Opened	Clicked-through	Active Users	Registered Users
TGroup1	153	61	22	14	6
TGroup2	288	88	20	14	6
CGroup	289	76	12	9	1

achieving the end goal. We will now examine each exit point sequentially to determine if the authority status of email senders affect the user outflow at each exit point.

4.1 Authority Effect on Email Open Rate

The first exit point we examine is when a user receives an email but ignores it completely, i.e., does not even open the email as illustrated in EP1 of Fig. 1. Thus, we are interested to see if the email sender affects the email open rates. TGroup1, TGroup2 and CGroup show an email open rate of 0.399, 0.306 and 0.263 respectively. One main observation is that TGroup1 outperforms CGroup by more than 51% in terms of email open rate, while TGroup2 manages a 15.8% improvement over the latter. This result indicates that the authority effect of the *superior* (HOD) has the greatest influence on email open rate, followed by that of the *domain expert* (event organizer), then the general account.

After examining the differences between the email open rates, we now determine if the differences are statistically significant. Using our collected data (Table 2), we conduct a two-sided Fisher’s Exact Test to determine if there is any significant difference for the following null hypotheses:

- H_{A0} : The *email open rate* of CGroup (*general account*) is equal to that of TGroup1 (*department head*).
- H_{B0} : The *email open rate* of CGroup (*general account*) is equal to that of TGroup2 (*event organizer*).

The resulting p -values are 0.0048 and 0.2691 for H_{A0} and H_{B0} respectively. At significance level $\alpha = 0.01$, we reject H_{A0} and fail to reject H_{B0} . There is moderate evidence to suggest that emails sent by the HODs are more likely to be opened by its recipients compared to emails from a general account. However, there is insufficient evidence to suggest that emails sent by the event organizer are statistically more effective than those sent by a general account.

4.2 Authority Effect on Email Open Time

In addition to the email open rate, we also investigate if users take less time to open emails from the HOD (TGroup1) or event organizer (TGroup2) compared to that from a general account (CGroup). Fig. 3 shows an Empirical Cumulative Distribution Function (ECDF) plot of the time taken by a user before he/she opens the email. The results show no observable difference among TGroup1, TGroup2 and CGroup in terms of the time taken to open the email.

To further support our observation that there is no difference in email open times, we test for statistical differences in the email open times by conducting a Two-sample Kolmogorov-Smirnov test on the following null hypotheses:

- H_{C0} : Distribution of *email open times* of CGroup (*general acct.*) is equal to that of TGroup1 (*dept. head*).
- H_{D0} : Distribution of *email open times* of CGroup (*general acct.*) is equal to that of TGroup2 (*event org.*).

The resulting p -values are 0.7724 and 0.3034 for H_{C0} and H_{D0} respectively. At significance level $\alpha = 0.01$, we fail to




Figure 3: Time to Email Open (left) and Click-thru (right)

reject H_{C0} and H_{D0} . Thus, there is insufficient evidence to show that emails sent by the HODs or event organizer are opened by its recipients in a different amount of time, compared to that sent by the general account.

This lack of difference in email open times could be explained by the mailbox checking routines of users. Users are likely to check their mailbox only at specific time intervals and upon checking their mailbox, they will look at the email preview (including sender name/address) before deciding whether to open the email. Thus, users in a group are unlikely to open an email faster than users in other groups as they would not have checked their emails. However, once they have checked their mailbox, users in TGroup1 and TGroup2 are more likely to proceed to open and read the email, as discussed previously in Sect. 4.1.

4.3 Authority Effect on Email Click Rate

After a user opens an email, the next possible exit point is when he/she simply closes the email after reading it, without clicking through to visit the event website, i.e., EP2 of Fig. 1. We now investigate if the authority status of email sender is effective in influencing the recipient to visit the event website. Our experiment shows that the email click-through rate of TGroup1, TGroup2 and CGroup are 0.144, 0.069 and 0.042 respectively. One main observation is that TGroup1 offered the best improvement with an increased email click-through rate of almost 3.5 times that of CGroup, while TGroup2 offered an improvement of 67% over CGroup. This result indicates that there is an authority effect on email click-through rates, with the authority type of *superior* (HOD) being most effective, followed by the *domain expert*, then the general email account.

Similar to the previous sections, we now determine if the differences in email click-through rate are statistically significant. Using Table 2, we conduct a two-sided Fisher’s Exact Test on the following null hypotheses:

- H_{E0} : The *email click-through rate* of CGroup (*general acct.*) is equal to that of TGroup1 (*dept. head*).
- H_{F0} : The *email click-through rate* of CGroup (*general acct.*) is equal to that of TGroup2 (*event org.*).

We obtained p -values of 0.0002 for H_{E0} and 0.1506 for H_{F0} . At significance level $\alpha = 0.01$, we reject H_{E0} and fail to reject H_{F0} . There is strong evidence to suggest that emails sent by HODs are more effective than those sent by the general account, in attracting the recipients’ attention to the extent that these recipients are interested in visiting the event website to find out more. There is insufficient evidence to suggest the same for emails sent by the event organizer.

4.4 Authority Effect on Email Click Time

In addition to the email click-through rate, we also investigate if the email sender affects the amount of time taken by

the recipient before he/she clicks through the email. Specifically, we are interested to see if a user takes a shorter time to click-through an email from the HOD (TGroup1) or event organizer (TGroup2) compared to that from a general account (CGroup). Fig. 3 shows that users who received emails from either their HOD (TGroup1) or the event organizer (TGroup2) are more likely to click-through in a shorter amount of time, compared to those received from the general account (CGroup). This result further reinforces our hypothesis that users receiving emails from their HODs and the event organizer are more likely to react to it and in a shorter time, compared to emails from a general account.

Having determined that authority status does affect email click-through time, we now examine if the difference is statistically significant by conducting a Two-sample Kolmogorov-Smirnov test on the following null hypotheses:

- H_{G0} : Distribution of *email click-through times* of CGroup (*general acct.*) is equal to TGroup1 (*dept. head*).
- H_{H0} : Distribution of *email click-through times* of CGroup (*general acct.*) is equal to TGroup2 (*event org.*).

The resulting p -values are <0.0001 for both H_{G0} and H_{H0} . At significance level $\alpha = 0.01$, we reject H_{G0} and H_{H0} . There is strong evidence to indicate that emails sent by the HODs or event organizer causes its recipient to click-through in a different amount of time, compared to emails sent by the general account.

4.5 Authority Effect on Website Activity

Even after receiving the email and clicking through, a user might simply visit the event website but leave without browsing any content, i.e., EP3 of Fig. 1. Thus, we now examine if the authority status of email sender is able to influence the recipients' behavior even on the event website. For this purpose, we define an active user as one who has interacted with any element on the event website, i.e., a mouse-click or mouse-over. Our experiment shows that the proportion of active users in TGroup1, TGroup2 and CGroup are 0.092, 0.049 and 0.031 respectively. This result once again shows the authority effect on user behavior in email campaigns, particularly that the authority type of *superior* (HOD) are almost two times more effective than the *domain expert* authority (event organizer). Both types of authority are also shown to be more effective than our control group without any authority (the general email account).

We now determine if the differences in website activity are statistically significant and use Table 2 to conduct a two-sided Fisher's Exact Test on the following null hypotheses:

- H_{I0} : Probability of *active users on the website* in CGroup (*general acct.*) is equal to TGroup1 (*dept. head*).
- H_{J0} : Probability of *active users on the website* in CGroup (*general acct.*) is equal to TGroup2 (*event org.*).

The resulting p -value are 0.0115 for H_{I0} and 0.2971 for H_{J0} . At significance level $\alpha = 0.05$, we reject H_{I0} and fail to reject H_{J0} . The significant difference in H_{I0} indicates that emails sent by the HODs effectively cause its recipient to be more active on the event website, compared to that sent by the general account. One significant improvement is that users in TGroup1 (9.2% active) are almost three times more likely to be active than users in CGroup (3.1% active), in terms of their browsing activities on the event website. While we do not observe any significant difference in H_{B0} , we notice that 4.9% of users in TGroup2 are active on the event website compared to only 3.1% of users in CGroup.

4.6 Authority Effect on Event Registration

The main goal of any email campaign is not only to entice users to read the email but ultimately to perform some specific action, such as registering for an event. We now further investigate if the authority status of email sender affects the conversion rate of users, i.e., do they register for the event or leave without registering, as shown in EP4 of Fig. 1. The registration rates of users in TGroup1, TGroup2, and CGroup are 0.039, 0.021 and 0.003 respectively, indicating that the authority types of *superior* (HOD) and *domain expert* (event organizer) are approximately eleven and six times more effective than that without authority, in terms of influencing users to register for the workshop.

Once again, we want to determine if the differences in event registration rates are statistically significant. Using Table 2, we conduct a two-sided Fisher's Exact Test on the following null hypotheses:

- H_{K0} : Probability of *event registration* in CGroup (*general acct.*) is equal to that of TGroup1 (*dept. head*).
- H_{L0} : Probability of *event registration* in CGroup (*general acct.*) is equal to that of TGroup2 (*event org.*).

We obtained p -values of 0.0080 and 0.0684 for H_{K0} and H_{L0} respectively. At significance level $\alpha = 0.01$, we reject H_{K0} and fail to reject H_{L0} . There is strong evidence to indicate a significant difference in H_{K0} , which shows that emails sent by HODs are more effective in getting users to register for the event, compared to emails sent by a general account. In particular, we observe a substantial and significant improvement of more than 11 times more converted users in TGroup1 (3.9% registration) compared to that of CGroup (0.3% registration). While we fail to reject H_{L0} , there is mild evidence (at significance level $\alpha = 0.1$) to suggest that emails sent by the event organizer (TGroup2) are six times more effective than emails sent by a general account (CGroup), in terms of converting users to register for the event.

5. CONCLUSION

In this paper, we designed and conducted an email campaign experiment to evaluate the authority effect of different types of email sender on user behavior in both email and website. We also proposed a systematic approach to analyze such email campaigns and discussed the four exit points where potential users might leave before completing a desired action or end-goal, such as registering for an event or purchasing an item. Other than experiment design, we developed a multi-platform experimentation system that manages and automates the key stages of controlled experiments, from user grouping randomization to data collection of user responses across multiple platforms, i.e., from emails to websites. Our findings show that the authority status of an email sender has a significant effect on the user response and emails sent by their department heads are the most effective in terms of: (i) email open rate; (ii) email click-through rate; (iii) email click-through time; (iv) proportion of active website users; and (v) event registration rate.

Acknowledgments. This research is supported by the Singapore National Research Foundation under its International Research Centre @ Singapore Funding Initiative and administered by the IDM Programme Office, Media Development Authority (MDA). The authors thank Feida Zhu, Adrian Vu, Agus Trisnaja Kwee and the reviewers for their useful comments.

6. REFERENCES

- [1] Customer.io, 2014. <http://customer.io/>.
- [2] Google analytics, 2014. <http://www.google.com/analytics>.
- [3] Mailchimp, 2014. <http://mailchimp.com/>.
- [4] S. Boulianne, C. A. Klofstad, and D. Basson. Sponsor prominence and responses patterns to an online survey. *International Journal of Public Opinion Research*, 23(1):79–87, 2011.
- [5] D. D. Castro, Z. Karnin, L. Lewin-Eytan, and Y. Maarek. You’ve got mail, and here is what you could do with it!: Analyzing and predicting actions on email messages. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining (WSDM’16)*, pages 307–316, 2016.
- [6] L. Chittenden and R. Rettie. An evaluation of e-mail marketing and factors affecting response. *Journal of Targeting, Measurement and Analysis for Marketing*, 11(3):203–217, 2003.
- [7] S. Cho, J. Huh, and R. J. Faber. The influence of sender trust and advertiser trust on multistage effects of viral advertising. *Journal of Advertising*, 43(1):100–114, 2014.
- [8] R. B. Cialdini. *Influence: Science and Practice*. Allyn & Bacon, 2001.
- [9] L. A. Dabbish, R. E. Kraut, S. Fussell, and S. Kiesler. Understanding email use: Predicting action on a message. In *Proc. of CHI’09*, pages 691–700, 2009.
- [10] W. Fan and Z. Yan. Factors affecting response rates of the web survey: A systematic review. *Computers in Human Behavior*, 26(2):132–139, 2010.
- [11] D. Fisher, A. J. Brush, E. Gleave, and M. A. Smith. Revisiting whittaker & sidner’s ”email overload” ten years later. In *Proceedings of the 20th conference on Computer supported cooperative work (CSCW’06)*, pages 309–312, 2006.
- [12] M. Grbovic, G. Halawi, Z. Karnin, and Y. Maarek. How many folders do you really need?: Classifying email into a handful of categories. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management (CIKM’14)*, pages 869–878, 2014.
- [13] C. Grevet, D. Choi, D. Kumar, and E. Gilbert. Overload is overloaded: email in the age of gmail. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI’14)*, pages 793–802, 2014.
- [14] R. Guadagno and R. B. Cialdini. Online persuasion and compliance: Social influence on the internet and beyond. *The social net: Human behavior in cyberspace*, pages 91–113, 2005.
- [15] N. Guéguen and C. Jacob. Solicitation by e-mail and solicitor’s status: A field study of social influence on the web. *CyberPsychology & Behavior*, 5(4):377–383, 2002.
- [16] R. Janke. Effects of mentioning the incentive prize in the email subject line on survey response. *Evidence Based Library and Information Practice*, 9(1), 2014.
- [17] M. D. Kaplowitz, F. Lupi, M. P. Couper, and L. Thorp. The effect of invitation design on web survey response rates. *Social Science Computer Review*, 30(3):339–349, 2012.
- [18] T. Karagiannis and M. Vojnovic. Behavioral profiles for advanced email features. In *Proceedings of the 18th international conference on World wide web (WWW’09)*, pages 711–720, 2009.
- [19] F. Kooti, L. M. Aiello, M. Grbovic, K. Lerman, and A. Mantrach. Evolution of conversations in the age of email overload. In *Proceedings of the 24th International Conference on World Wide Web (WWW’15)*, pages 603–613, 2015.
- [20] K. H. Lim, B. Jiang, E.-P. Lim, and P. Achananuparp. Do you know the speaker? An online experiment with authority messages on event websites. In *Proceedings of the 23rd International Conference on World Wide Web Companion (WWW’14)*, pages 1247–1252, 2014.
- [21] K. H. Lim, E.-P. Lim, P. Achananuparp, A. Vu, A. T. Kwee, and F. Zhu. LASER: A Living AnalyticS Experimentation system for large-scale online controlled experiments. In *Proceedings of the 23rd International Conference on World Wide Web Companion (WWW’14)*, pages 71–74, 2014.
- [22] B.-W. On, E.-P. Lim, J. Jiang, and L.-N. Teow. Engagingness and responsiveness behavior models on the enron email network and its application to email reply order prediction. In *The Influence of Technology on Social Network Analysis and Mining*, volume 6 of *LNCS*, pages 227–253. Springer Vienna, 2013.
- [23] J. E. Phelps, R. Lewis, L. Mobilio, D. Perry, and N. Raman. Viral marketing or electronic word-of-mouth advertising: Examining consumer responses and motivations to pass along email. *Journal of Advertising Research*, 44(4):333–348, 2004.
- [24] Radicati Group, Inc. Email statistics report 2015-2019, 2015. <http://www.radicati.com/wp/wp-content/uploads/2015/07/Email-Market-2015-2019-Executive-Summary.pdf>.
- [25] Salesforce.com. 25 mind blowing email marketing stats, 2013. <https://www.salesforce.com/blog/2013/07/email-marketing-stats.html>.
- [26] W. Shu and C. Y. Cheng. How to improve consumer attitudes toward using credit cards online: An experimental study. *Electronic Commerce Research and Applications*, 11(4):335–345, 2012.
- [27] J. C. Tang, T. Matthews, J. Cerruti, S. Dill, E. Wilcox, J. Schoudt, and H. Badenes. Global differences in attributes of email usage. In *Proceedings of the 2009 International Workshop on Intercultural Collaboration (IWIC’09)*, pages 239–258, 2009.
- [28] J. R. Tyler and J. C. Tang. When can i expect an email response? a study of rhythms in email usage. In *Proceedings of the Eighth European Conference on Computer Supported Cooperative Work (ECSCW’03)*, pages 239–258, 2003.
- [29] S. Whittaker and C. Sidner. Email overload: Exploring personal information management of email. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI’09)*, pages 276–283, 2009.
- [30] S. Yoo, Y. Yang, F. Lin, and I.-C. Moon. Mining social networks for personalized email prioritization. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD’09)*, pages 967–976, 2009.