



Publicly Accessible Penn Dissertations

---

1-1-2016

# Essays on Social Media and Digital Marketing

Jing Peng

University of Pennsylvania, pengjing07@gmail.com

Follow this and additional works at: <http://repository.upenn.edu/edissertations>

 Part of the [Advertising and Promotion Management Commons](#), and the [Marketing Commons](#)

---

## Recommended Citation

Peng, Jing, "Essays on Social Media and Digital Marketing" (2016). *Publicly Accessible Penn Dissertations*. 1942.  
<http://repository.upenn.edu/edissertations/1942>

This paper is posted at ScholarlyCommons. <http://repository.upenn.edu/edissertations/1942>  
For more information, please contact [libraryrepository@pobox.upenn.edu](mailto:libraryrepository@pobox.upenn.edu).

---

# Essays on Social Media and Digital Marketing

## **Abstract**

Digital technology is rapidly reshaping the way how brands interact with consumers. More and more marketers are shifting their focus from traditional marketing channels (e.g., TV) to digital channels (e.g., social media platforms). Effective targeting is key to successful social media and digital marketing campaigns. This dissertation seeks to shed light on who and how to target on social media platforms.

The first chapter aims to provide insights on how to target customers who are connected to each other on social media platforms. We investigate how the network embeddedness (i.e., number of common followees, common followers, and common mutual followers) between two users impacts information diffusion from one (sender) to another (receiver). By analyzing the sharing of sponsored ads on Digg and brand-authored tweets on Twitter, we find that the effect of embeddedness in directed networks varies across different types of “neighbors”. A receiver is more likely to share content from a sender if they share more common followees. A receiver is also more likely to share content if she shares more common followers and common mutual followers with the sender. However, this effect is moderated by the novelty of information.

The second chapter strives to understand what affects paid endorsers’ participation and effectiveness in social advertising campaigns. We conduct a field experiment with an invitation design in which we manipulate both incentives and a soft eligibility requirement to participate in the campaign. There are three main findings from our analysis. (1) Payments higher than the average reward a potential endorser received in the past (gains) do not increase participation, whereas lower payments (losses) decrease participation. Neither gains nor losses compared to past reward affect performance. (2) Potential endorsers who are more likely to participate tend to be less effective. (3) Which characteristics are associated with effectiveness depends on whether success is measured in likes, comments, or retweets.

For marketing managers, our findings provide insights on how to target customers in a directed network at a micro level and how to improve social advertising campaigns by better targeting and incenting potential endorsers.

## **Degree Type**

Dissertation

## **Degree Name**

Doctor of Philosophy (PhD)

## **Graduate Group**

Operations & Information Management

## **First Advisor**

Kartik Hosanagar

## **Second Advisor**

Christophe Van den Bulte

---

**Keywords**

content sharing, digital marketing, paid endorsement, social media marketing, targeting

**Subject Categories**

Advertising and Promotion Management | Marketing

ESSAYS ON SOCIAL MEDIA AND DIGITAL MARKETING

Jing Peng

A DISSERTATION

in

Operations, Information and Decisions

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the  
Degree of Doctor of Philosophy

2016

Supervisor of Dissertation

Co-Supervisor of Dissertation

---

Kartik Hosanagar  
Professor of Operations, Information and Decisions

---

Christophe Van den Bulte  
Professor of Marketing

Graduate Group Chairperson

---

Catherine Schrand, John C. Hower Professor; Professor of Accounting

Dissertation Committee:

Lorin Hitt, Professor of Operations, Information and Decisions  
Raghuram Iyengar, Associate Professor of Marketing

## DEDICATION

*This dissertation is dedicated to my parents Huagui Peng and Qinmei Zhang. Thank you isn't enough for your unwavering support and encouragement throughout this long journey.*

## ACKNOWLEDGMENT

Foremost, I would like to express my sincere gratitude to my advisor Prof. Kartik Hosanagar for his open-mindedness and continuous support on my research. He is extremely patient and always encourage me to pursue whatever I am interested in. I am deeply grateful to my co-advisor Prof. Christophe Van den Bulte for his enthusiasm and immense knowledge. He is always there for me. I can't remember how many times I walked into his office and talked with him for hours without making appointments in advance.

My sincere thanks also goes to the rest of my thesis committee: Prof. Lorin Hitt and Prof. Raghuram Iyengar, for their insightful and constructive comments, which incited me to sharpen my research from various perspectives. I would like to especially thank my co-author Prof. Ashish Agarwal for his tremendous help in my research.

I would also like to take this opportunity to thank my friends and the fellow Wharton PhD students for their invaluable help.

Last but not the least, I appreciate the financial support from the Operations, Information and Decisions Department, the Baker Center PhD Research Grant, the Mack Institute Research Fellowship, the Penn Lauder CIBER PhD Grant, and the President Gutmann's Leadership Award that funded parts of the research discussed in this dissertation.

## ABSTRACT

### ESSAYS ON SOCIAL MEDIA AND DIGITAL MARKETING

Jing Peng

Kartik Hosanagar

Christophe Van den Bulte

Digital technology is rapidly reshaping the way how brands interact with consumers. More and more marketers are shifting their focus from traditional marketing channels (e.g., TV) to digital channels (e.g., social media platforms). Effective targeting is key to successful social media and digital marketing campaigns. This dissertation seeks to shed light on who and how to target on social media platforms.

The first chapter aims to provide insights on how to target customers who are connected to each other on social media platforms. We investigate how the network embeddedness (i.e., number of common followees, common followers, and common mutual followers) between two users impacts information diffusion from one (sender) to another (receiver). By analyzing the sharing of sponsored ads on Digg and brand-authored tweets on Twitter, we find that the effect of embeddedness in directed networks varies across different types of “neighbors”. A receiver is more likely to share content from a sender if they share more common followees. A receiver is also more likely to share content if she shares more common followers and common mutual followers with the sender. However, this effect is moderated by the novelty of information.

The second chapter strives to understand what affects paid endorsers’ participation and effectiveness in social advertising campaigns. We conduct a field experiment with an invitation design in which we manipulate both incentives and a soft eligibility requirement to participate in the campaign. There are three main findings from our analysis. (1) Payments higher than the average reward a potential endorser received in the past (gains) do not increase participation, whereas lower payments (losses) decrease participation. Neither gains nor losses compared to past reward

affect performance. (2) Potential endorsers who are more likely to participate tend to be less effective. (3) Which characteristics are associated with effectiveness depends on whether success is measured in likes, comments, or retweets.

For marketing managers, our findings provide insights on how to target customers in a directed network at a micro level and how to improve social advertising campaigns by better targeting and incenting potential endorsers.

TABLE OF CONTENTS

**DEDICATION ..... II**

**ACKNOWLEDGMENT ..... III**

**ABSTRACT ..... IV**

**LIST OF TABLES ..... VIII**

**LIST OF ILLUSTRATIONS ..... IX**

**1. NETWORK EMBEDDEDNESS AND CONTENT SHARING ON SOCIAL MEDIA PLATFORMS..... 1**

1.1 Introduction ..... 1

1.2 Related Literature..... 5

1.3 Theoretical Background and Hypotheses ..... 8

    1.3.1 Common Followees ..... 10

    1.3.2 Common Followers ..... 10

    1.3.3 Common Mutual Followers ..... 11

1.4 Model..... 13

    1.4.1 Dyadic Hazard..... 14

    1.4.2 Spontaneous Sharing ..... 15

    1.4.3 Model Estimation..... 16

    1.4.4 Identification ..... 18

1.5 Data ..... 19

1.6 Results ..... 25

    1.6.1 Main Results ..... 25

    1.6.2 Robustness Checks ..... 27

    1.6.3 Generalizability to Other Social Networks ..... 30

1.7 Discussion & Conclusion..... 32

**2. PARTICIPATION VS. EFFECTIVENESS OF PAID ENDORSERS IN SOCIAL ADVERTISING CAMPAIGNS: A FIELD EXPERIMENT ..... 36**

2.1 Introduction ..... 36

2.2 Theoretical Background..... 39

2.2.1 Participation .....	39
2.2.2 Effectiveness.....	41
<b>2.3 Field Experiment .....</b>	<b>43</b>
2.3.1 Research Setting.....	43
2.3.2 Experiment Design.....	44
<b>2.4 Data .....</b>	<b>47</b>
2.4.1 Descriptive Statistics .....	47
2.4.2 Model-Free Analysis of Manipulation Effects.....	51
<b>2.5 Model.....</b>	<b>52</b>
2.5.1 Sample Selection Model with Correlated Random Effects .....	52
2.5.2 Connections with Existing Models .....	55
2.5.3 Relative Partial Effects on Potential and Actual Outcome .....	56
<b>2.6 Results .....</b>	<b>57</b>
2.6.1 Selection of Incentive Variables.....	57
2.6.2 Main Results .....	58
2.6.3 Robustness .....	62
<b>2.7 Implications for Program Design .....</b>	<b>64</b>
2.7.1 Influencing Endorsers by Redesigning Tasks.....	64
2.7.2 Boosting Potential vs. Actual Engagements by Targeting .....	66
<b>2.8 Conclusions.....</b>	<b>68</b>
 <b>APPENDIX .....</b>	 <b>70</b>
<b>Appendix 1.1: Simulation Studies .....</b>	<b>70</b>
<b>Appendix 1.2: Complete Results .....</b>	<b>73</b>
<b>Appendix 1.3 Diffusion Graphs .....</b>	<b>74</b>
<b>Appendix 1.4 Additional Results .....</b>	<b>76</b>
<b>Appendix 2.1 Likelihood and Parameter Estimation .....</b>	<b>79</b>
<b>Appendix 2.2 Mean Actual Outcome.....</b>	<b>81</b>
<b>Appendix 2.3 Relative Partial Effects.....</b>	<b>82</b>
<b>Appendix 2.4 Prediction of Participation and Effectiveness .....</b>	<b>84</b>
<b>Appendix 2.5 Results Using Alternative Incentives .....</b>	<b>85</b>
<b>Appendix 2.6 Additional Robustness Checks .....</b>	<b>89</b>
 <b>BIBLIOGRAPHY .....</b>	 <b>92</b>

## LIST OF TABLES

### Chapter 1

Table 1.1 Glossary .....	2
Table 1.2 Literature on the Role of Network Characteristics on User Actions.....	7
Table 1.3 Drivers Associated with the Three Embeddedness Metrics .....	12
Table 1.4 Summary Statistics .....	22
Table 1.5 Descriptions of Independent Variables .....	23
Table 1.6 Key Statistics of Main Variables.....	23
Table 1.7 Correlation among Dyadic Network Characteristics .....	24
Table 1.8 Parameters Estimates of Different Model Specifications.....	25
Table 1.9 Parameters Estimates from Different Random/Fixed/Mixed Effects Models.....	28
Table 1.10 Parameters Estimates on the Two Subsets.....	29
Table 1.11 Parameter Estimates on Twitter Dataset .....	31

### Chapter 2

Table 2.1 Experimental Design .....	46
Table 2.2 Experiment Statistics .....	47
Table 2.3 Participation and Engagement Statistics of Tasks.....	48
Table 2.4 Distribution of Engagements Generated by Individual Endorsers .....	48
Table 2.5 Description of Independent Variables .....	49
Table 2.6 Key Statistics on Independent Variables .....	50
Table 2.7 Correlation between Independent Variables.....	50
Table 2.8 Effects of Manipulated Variables .....	52
Table 2.9 Models for Panel Count Data with Sample Selection .....	56
Table 2.10 Selection of Incentive Variables.....	58
Table 2.11 Parameter Estimates for Different Types of Engagements (Log GainLoss) .....	59
Table 2.12 Distribution of Endorsers.....	65
Table 2.13 (Relative) Partial Effects on Participation, Potential and Actual Effectiveness.....	67

### Appendix 1

Table A1.1 Relative Errors of the Collective Cause Model.....	72
Table A1.2 Complete Parameter Estimates for Models in Table 1.8.....	73
Table A1.3 Parameter Estimates on the Digg Dataset (including inactive users) .....	76
Table A1.4 Summary Statistics for Twitter Dataset .....	76
Table A1.5 Descriptions of Independent Variables for Twitter Dataset .....	77
Table A1.6 Key Statistics of Main Variables for Twitter Dataset.....	77
Table A1.7 Correlation among Dyadic Network Characteristics for Twitter Dataset .....	78
Table A1.8 Complete Results on the Twitter Dataset.....	78

### Appendix 2

Table A2.1 Parameter Estimates for Different Types of Engagements (Pay Rate).....	85
Table A2.2 Parameter Estimates for Different Types of Engagements (Linear Reward) .....	86
Table A2.3 Parameter Estimates for Different Types of Engagements (Linear GainLoss) .....	87
Table A2.4 Parameter Estimates for Different Types of Engagements (Log Reward) .....	88
Table A2.5 Effect of Eligibility on Effort Level .....	89
Table A2.6 Robustness to Model Complexity .....	90
Table A2.7 Robustness to Potential Outliers .....	91

## LIST OF ILLUSTRATIONS

### **Chapter 2**

Figure 2.1 The Incentive Curves at Low and High Pay Rates..... 45

### **Appendix 1**

Figure A1.1 Kaplan-Meier Survival Curve for Digg Ads..... 74

Figure A1.2 Sharing Graphs for Ads..... 75

# 1. Network Embeddedness and Content Sharing on Social Media Platforms

## 1.1 Introduction

Social media platforms are a popular medium for firms to reach out to customers (Schweidel and Moe 2014; Stephen and Toubia 2010). A recent survey suggests that three quarters of advertisers had used social media for advertising, and 64% of them planned to increase their social advertising budgets (Nielsen 2013). One likely reason for the growing emphasis on social advertising is the promise that users who engage with the ad content might spread information about new products to their social network connections (Aral and Walker 2011; Aral and Walker 2012; Aral and Walker 2014; De Bruyn and Lilien 2008; Leskovec et al. 2007; Trusov et al. 2010).

A primary requirement for the propagation of content in a social network is that receivers in turn share the information that they obtain from their sender/s. However, empirical evidence for such information cascades is limited (Goel et al. 2012). For instance, the average number of retweets per tweet is often less than 20.<sup>1</sup> Thus, it is important to understand the underlying drivers for the propensity of a receiver to share information. Such insights will also be useful for marketers to improve their targeting of customers within social networks (Kempe et al. 2003; Richardson and Domingos 2002; Watts and Dodds 2007).

Extant work indicates that one important driver of the propensity of a receiver to share information is the embeddedness or network overlap between a dyad, i.e., a sender-receiver pair (Aral and Walker 2014). Network embeddedness or network overlap<sup>2</sup> is a shared characteristic between users in a network and has been associated with effective knowledge transfer between individuals (Reagans and McEvily 2003), extent of information sharing among users (Aral and Van Alstyne

---

<sup>1</sup> Social engagement benchmark report (salesforce 2015): <https://www.marketingcloud.com/resource-center/digital-marketing/benchmark-2014/social-engagement-tw/>

<sup>2</sup> We use embeddedness and network overlap interchangeably in this paper

2011) and adoption of applications by users (Aral and Walker 2014). In the context of firms, network embeddedness has been associated with trust between firms (Uzzi 1997) and their economic actions (Granovetter 1985).

Table 1.1 Glossary

Glossary	Description
<b>Connections</b>	
Friend	A user mutually connected with the focal user (undirected networks)
Followee	A user followed by the focal user (directed networks)
Follower	A user following the focal user (directed networks)
Mutual follower	A user following and followed by the focal user (directed networks)
<b>Embeddedness</b>	
Common friend	A user mutually connected to both the sender and the receiver (undirected networks)
Common followee	A user followed by both the sender and the receiver (directed networks)
Common follower	A user following both the sender and the receiver (directed networks)
Common mutual follower	A user following and followed by both the sender and the receiver (directed networks)
<b>Others</b>	
Share	Digg an ad or retweet a tweet
Feed	information notifying a user about the sharing activity of one's followees
Co-senders	The set of followees of the focal user who have already shared the ad/tweet

Network embeddedness is broadly defined as the number of common neighbors between two users (Easley and Kleinberg 2010). Its operationalization depends on whether the network is directed or not. In undirected networks (e.g., Facebook), embeddedness or network overlap simply means the number of common friends between two users. In directed networks (e.g., Twitter and Digg<sup>3</sup>), by interpreting a neighbor as a followee (outgoing link), follower (incoming link) or mutual follower (bidirectional link), embeddedness can be characterized by three different metrics: the numbers of common followees, common followers, and common mutual followers. Table 1.1 summarizes the definitions of these terms. The distinction between followers and followees is important. In directed networks like Twitter and Weibo, one can follow a user without consent from the user. Followees of a focal user thereby represent the set of users whose activities are of interest to the focal user,

<sup>3</sup> Digg maintained an internal directed network before August 2012, but now it uses the external social networks of users on Twitter and Facebook instead.

whereas the followers represent the set of users who are interested in the focal user's activity. Mutual follower (a bidirectional link) cannot be established unless users have mutual interest.

The nature of overlap in the network connections between two users can reveal the motivation to share content. For example, followees of a user can have a high persuasive influence on the user (Haenlein 2013; Hall and Valente 2007) and can represent user's interests and expertise. Furthermore, people tend to share content that signals their expertise (Packard & Wooten, 2013). Thus, more common followees between a sender-receiver pair may suggest similar expertise and higher propensity for a receiver to share the content obtained from the sender. Likewise, more common followers between the sender and the receiver may suggest that their followers share a similar taste. In this case, a receiver may consider content to be more suitable for her audience and may have a higher propensity to share content. Additionally, a higher number of common mutual followers may represent higher trust (Burt 2001; Granovetter 1973) and social bonding (Alexandrov et al. 2013; Ho and Dempsey 2010; Travis 2002; Wiatrowski et al. 1981) and may also increase the propensity of a receiver to share information. Finally, as the activities of users on social media platform are visible to others, factors such as uniqueness of content can play a role. It is well documented that users (consumers) have a strong desire for uniqueness and sharing novel content can satisfy such a need (Alexandrov et al. 2013; Cheema and Kaikati 2010; Ho and Dempsey 2010; Lovett et al. 2013). Thus, if the information to be communicated is not novel, a receiver will be less likely to do so.

The purpose of this article is to assess the impact of embeddedness on the level of content sharing in directed networks. We do so using a micro-level model for content sharing within sender-receiver dyads. Our work complements extant work on the role of influential users on product adoption (Iyengar et al. 2011; Trusov et al. 2010) and information diffusion (Susarla et al. 2012; Yoganasimhan 2012). Other studies have described a user's propensity to adopt a product and share related information based on unitary attributes of adopters such as their demographic and behavioral characteristics (Bapna and Umyarov 2015; Haenlein 2013; Iyengar et al. 2011; Katona

et al. 2011; Nitzan and Libai 2011; Rand and Rust 2011). Some other studies have considered shared characteristics of a sender and a receiver but largely in undirected networks (Aral and Walker 2014) where there is a single metric for the overlap among users, i.e., the number of common friends. Finally, in the case of directed networks, to the best of our knowledge, only the effect of reciprocity in the connections between a sender and a receiver has been considered (Shi et al. 2014).

A dyadic level study of content sharing imposes stringent requirements on the data: the availability of users' profile information, social graph information, and time-stamped, highly granular, individual-level information about sharing activities. In order to meet these requirements, we constructed a dataset, which represents sharing of sponsored ads on Digg in a month long period in 2012. We corroborate our results using a second dataset that captures sharing of tweets posted by Fortune 500 companies on Twitter in a month long period in 2015. At the time the data were collected, both websites maintained a directed social network, allowing users to follow others to keep themselves informed about their activities.

A dyadic level study introduces a methodological challenge as well: multiple senders may share the same content with a focal user and the lack of information regarding the contribution of each sender makes it difficult to identify the impact of dyadic characteristics on receiver's sharing propensity. For example, in the dataset from Digg, 32% of receivers who decided to share had multiple co-senders. In response to this problem, we propose a novel proportional hazards model that allows an event to have more than one cause. The proposed model can identify the contribution of each co-sender based on her characteristics and has broader application in studies of diffusion in networks.

We emerge from the analysis with three key findings. First, we establish that embeddedness plays a role in information sharing in directed networks. That is, the propensity of receiver to share information depends on all three measures of embeddedness (i.e. common followees, common followers and common mutual followers). Second, the effect of embeddedness on content sharing

varies across the three metrics suggesting that they may have differing underlying drivers. Third, the effects of common followers and common mutual followers are moderated by the novelty of content. Their effects are positive only when the information is relatively novel (i.e., not shared by many others). When many others have shared the content, the positive effects may decrease and may even become negative, likely due to users' need for uniqueness. This finding suggests a boundary condition for the positive impact of embeddedness found in previous work.

The rest of the paper is organized as follows. We begin with a discussion of related literature and propose specific hypotheses about the impact of the three embeddedness metrics on content sharing. Then, we describe the proposed model and the dataset from Digg that we use in the application. Next, we discuss the results of model estimation and several robustness checks including validation of our results with the dataset from Twitter. Finally, we conclude with theoretical and managerial implications.

## **1.2 Related Literature**

Our work relates to the broad literature on the role of network characteristics on user actions and outcomes in a social network. These include studies of information sharing (Shi et al. 2014; Susarla et al. 2012; Yoganarasimhan 2012), product adoption (Aral and Walker 2014; Bapna and Umyarov 2015; Iyengar et al. 2011; Katona et al. 2011), and customer churn (Haenlein 2013; Nitzan and Libai 2011).

Some studies have investigated the role of unitary network characteristics of the sender on the overall extent of adoption in the network. For example, Yoganarasimhan (2012) studies how the size and structure of the local network of a user affect product diffusion in undirected networks. The specific network characteristics investigated include the numbers of first- and second-degree friends, the clustering coefficient and the centrality of the user. Susarla et al. (2012) conduct a similar analysis but include both undirected (friendship) and directed (subscription) networks on Youtube. Bakshy et al. (2011) determine the user influence based on the cascade size associated

with a user's extended network. While these studies consider the effect of sender's local and extended network on their effectiveness in spreading product adoption behavior, they do not consider an individual receiver's propensity to adopt these products.

Others have investigated the role of unitary network characteristics of the receiver on her individual adoption behavior. For instance, Iyengar et al. (2011) consider the impact of user characteristics such as opinion leadership (captured by the number of ties and self-reported measures) on the adoption of prescription drugs. Katona et al. (2011) investigate the effect of a receiver's network characteristics on their adoption or registration at a site. Similarly, Bapna and Umyarov (2015) consider the effect of the receiver's network size on her propensity to subscribe to a music site. Hu and Van den Bulte (2014) focus on status characteristics and network centrality. Rand and Rust (2011) evaluate the role of local network on the adoption behavior using an agent based model. Nitzan and Libai (2011) and Heinlein (2013) investigate the role of the network neighbors' churn behavior on the retention behavior of an individual. However, none of the above studies considers the impact of shared network characteristics between the receiver and the sender on the former's adoption behavior.

Some recent studies do focus on the role of shared characteristics on a focal user's actions albeit in undirected networks. For example, Centola (2010) shows that users are more likely to adopt when they receive social reinforcement from multiple neighbors and, as a result, the behavior spreads more in a clustered network than a random network. While a clustered network can represent higher network overlap with neighbors, this overlap is artificially created in the experiment and does not directly capture the shared characteristics between two users. Aral and Van Alstyne (2011) investigate the role of embeddedness on the sender's incentive to share information with a particular receiver but not the receiver's propensity to, in turn, share the content with all her followers. Aral and Walker (2014) examine the effect of network embeddedness more directly and find that it has a positive effect on the adoption of an application on Facebook (an undirected

network). Finally, while Shi et al. (2014) study information sharing in a directed network, they primarily focus on the role of reciprocity between senders and receivers.

Table 1.2 Literature on the Role of Network Characteristics on User Actions

Study	Network Characteristics	Network Type	Context and User Actions
	<b>Dyadic network characteristics</b>		
<b>Present study</b>	Three network embeddedness metrics between dyads	Directed	Online content sharing
Aral and Walker (2014)	Network embeddedness and interaction intensity between dyads	Undirected	Facebook app adoption
Aral and Van Alstyne (2011)	Network embeddedness between dyads	Undirected	Information sharing by sender with individual receiver
Shi et al (2014)	Reciprocity between dyads	Directed	Online content sharing
	<b>Unitary network characteristics</b>		
Yoganarasimhan (2012)	Network characteristics of sender	Undirected	Diffusion of Youtube videos and related information
Susarla et al. (2012)	Network characteristics of sender	Directed and Undirected	Diffusion of Youtube videos and related information
Katona et al. (2011)	Network characteristics of receiver	Undirected	Registration (Adoption) of social networking site
Bapna and Umyarov (2015)	Network size of receivers	Undirected	Subscription (Adoption) of Last.fm
Iyengar et al. (2011)	Opinion leadership of receiver	Directed	Adoption of prescription drug
Hu and Van den Bulte (2014)	Social status of receiver	Directed and Undirected	Adoption of site-directed mutagenesis kits
Bakshy et al. (2011)	Network characteristics of sender	Directed	Information diffusion
Rand and Rust (2011)	Network characteristics of receiver	Undirected	Adoption behavior using an agent based model
Irit and Libai (2011)	Churn of behavior of neighbors	Undirected	Churn behavior of receiver
Hanlein (2013)	Churn behavior of ingoing and outgoing connections of receiver	Directed	Churn behavior of receiver
	<b>Overall network structure</b>		
Centola (2010)	Overall structure of network (clustered vs. random)	Undirected	Registration (Adoption) of online health forum

In summary, there is clearly much interest in understanding how users' network characteristics affect product diffusion and information sharing in networks. While previous work has focused on either aggregate network measures or unitary characteristics of senders and receivers, an emerging stream of work has started to highlight the role of such dyadic attributes as network embeddedness. This literature, to the best of our knowledge, has considered undirected networks. In this paper, we fill the gap and evaluate how network embeddedness affects information sharing in directed networks. Table 1.2 provides a summary of existing literature.

### 1.3 Theoretical Background and Hypotheses

Consumers typically share content to satisfy multiple goals. Users may share content with others in a social network for the purpose of impression management (Berger 2014; Toubia and Stephen 2013). Further, factors such as trustworthiness of a sender (Burt 2001; Granovetter 1973) may play a role in a user's propensity to share any content received from a sender. Users may have additional motives as well to share content such as social bonding (Alexandrov et al. 2013; Ho and Dempsey 2010; Travis 2002; Wiatrowski et al. 1981) and the need for uniqueness (Cheema and Kaikati 2010; Grier and Deshpandé 2001; Ho and Dempsey 2010; Lovett et al. 2013; Snyder and Fromkin 1980). Next, we outline these motivations in more detail and how they relate to our main construct of network embeddedness.

*Impression Management.* Users share content to shape others' impression about them. On social media platforms, users' activities are publicly visible to others. Such visibility of individual activities makes social media platforms an ideal place to create an impression and enhance their social status (Alexandrov et al. 2013; Grier and Deshpandé 2001; Lovett et al. 2013; Toubia and Stephen 2013).

Users may try to impress others by communicating specific identities (Berger 2014). For instance, people share topics or ideas that signal that they have certain characteristics, knowledge base or expertise (Packard & Wooten, 2013). Further, content sharing is a social exchange process (Aral and Van Alstyne 2011). To increase social acceptance or social recognition, users may selectively share information of interest to their audience (Aral and Van Alstyne 2011; Wu et al. 2004), as sharing information perceived to be unsound or irrelevant can hurt their reputation (Barasch and Berger 2014; Bock et al. 2005).

*Trust.* Trust is a key determinant of social information exchange process (Burt 2001; Granovetter 1973). The trust of users on the source (i.e., senders) can alleviate the receivers' concern on the quality of the content and hence increase the probability of sharing (Camarero and San José 2011).

The trust between two users often increases with common mutual connections (bandwidth) between them (Aral and Van Alstyne 2011; Burt 2001).

*Social bonding.* Social control theory suggests that people have a need to bond with others and maintain relationships (Travis 2002; Wiatrowski et al. 1981). Social bonding is also referred to as “need to belong” (Alexandrov et al. 2013; Ho and Dempsey 2010). The formation of a bond between individual and a group requires frequent interactions with others in the group (Alexandrov et al. 2013). On social media platforms, as the user actions are visible, one way to interact with others is to further share the content shared by others. The closer two users are, the stronger obligation they may have in sharing content shared by each other.

*Need for uniqueness.* The theory on self-presentation suggests that users are intrinsically motivated to achieve uniqueness (Tajfel and Turner 1979; Turner and Oakes 1986) and being overly similar to others induces negative emotions (Snyder and Fromkin 1980). This desire to express uniqueness is stronger for publicly consumed products than privately consumed products (Cheema and Kaikati 2010). Moreover, the need for uniqueness is stronger in online interactions than offline interactions and leads to higher word of mouth for differentiated brands (Lovett et al. 2013). Need for uniqueness has also been observed for other user generated content such as reviews (Ludford et al. 2004) and photographs (Zeng and Wei 2013). Thus, in order to establish a unique identity on social media platforms, a user may resist sharing content that have already been shared by many others.

In the case of content received from a sender, we posit that the characteristics shared between the user and the sender are an important contextual feature that can moderate how likely a user will satisfy one or more of the above mentioned goals and, thereby, influence their propensity to share content. We use network overlap or embeddedness between users in a social network to operationalize the shared characteristics. Next, we discuss our hypotheses on how the three metrics of embeddedness can impact content sharing.

### 1.3.1 Common Followees

In a directed social network, people follow others to keep themselves informed about their activities. Followees of a user can have a high persuasive influence on the user (Haenlein 2013; Hall and Valente 2007). Thus, the composition of one's followees largely reflects her topical interest or taste. In addition to taste, the composition of one's followees may also reflect her expertise, as people may selectively follow others with similar expertise. In order to signal online identities and create an impression, users tend to share content falling into their area of expertise or interest (Berger 2014; Packard & Wooten, 2013). This is likely irrespective of the type of content, including popular content. Therefore, the more common followees two users have, the more likely they have similar expertise and taste due to homophily, and the more likely they will share the content shared by each other. So we posit the following:

*H1: The propensity of a receiver to share a piece of content from a sender is positively associated with the number of common followees between the sender and the receiver.*

### 1.3.2 Common Followers

The composition of one's followers represents the taste of her audience. To establish a good impression, the taste of audience is an important factor that users are likely to consider while sharing content (Berger 2014). The more common followers two users have, the more similar audience they have, and the more likely they will make similar decisions on whether or not to share a piece of content to their followers to create an impression.

On the other hand, an alternative driver that may lower the propensity of a receiver to share content obtained from a sender with whom the receiver has a lot of common followers is the need for uniqueness. Sharing redundant (i.e., duplicated) content that has already been seen by their followers from other sources can harm the perception of the receiver as a unique source of information. Thus, novelty of content can play a role in moderating the impact of common followers on the propensity of a receiver to share content. Less popular or novel information is more valuable

due to its scarcity (Aral and Van Alstyne 2011). When the content is not as popular yet, the novelty of the content will make it relatively easier for a receiver to distinguish herself from others. In such a case, the sharing decision of the receiver should be primarily driven by impression management rather than by her need for uniqueness (as it is being satisfied by sharing novel content). When the content is popular, the need for uniqueness may be strong enough to outweigh impression management. Following these arguments, we propose the following hypotheses.

*H2: The propensity of a receiver to share a piece of content from a sender is positively associated with the number of common followers between the sender and the receiver.*

*H3: The positive effect of common followers on the receiver's propensity to share content from sender decreases with the popularity of the information.*

### 1.3.3 Common Mutual Followers

The number of common mutual followers characterizes the mutual accessibility of two users through third-parties, which may be the most appropriate counterpart to the embeddedness defined in undirected networks. According to the bandwidth hypothesis (Aral and Van Alstyne 2011; Burt 2001), the existence of common mutual connections expands the bandwidth of communication among users and makes their evaluation of each other more accurate. Therefore, the level of trust between two users should increase with the number of common mutual followers. In addition, the more common mutual followers two users have, the more likely they belong to the same social group, and the more likely they feel obligated to propagate content shared by each other in order to maintain a strong social bond. Both drivers on trust and social bonding suggest that the number of common mutual followers should have positive effect on content sharing. More common mutual followers may also suggest a common taste for audience. Finally, more common mutual followers suggests higher similarity in taste and expertise due to homophily even after accounting for the effect of other embeddedness metrics. This would further increase the receiver's propensity to share content.

However, a user's need for uniqueness can lower her propensity to share content from a sender with whom she shares mutual common followers. Similar to our earlier reasoning for the effect of common followers on content sharing, when the content to be shared is popular, a receiver with a large number of common mutual followers with a sender may resist doing so to avoid excessive similarity with the sender, as well as with other members in the same social group. However, when the content is relatively novel, the need for uniqueness is already satisfied and the receiver would have a higher propensity to share content due to high number of common mutual followers. We summarize the expected effects of common mutual followers in H4 and H5.

*H4: The propensity of a receiver to share a piece of content from a sender is positively associated with the number of common mutual followers between the sender and the receiver.*

*H5: The positive effect of common mutual followers on the receiver's propensity to share content from sender decreases with the popularity of the information.*

Table 1.3 summarizes the drivers associated with the three embeddedness metrics in directed networks. Note that the need for uniqueness as a driver should only come into play when there is an audience. Thus, the need for uniqueness is unlikely to drive the effect of common followees, as followees represent sources rather than the audience of a focal user. That different drivers are associated with the three metrics illustrates the nuanced role of embeddedness on information sharing in directed networks.

Table 1.3 Drivers Associated with the Three Embeddedness Metrics

<b>Embeddedness Metric</b>	<b>Positive Driver</b>	<b>Negative Driver</b>
Common followees	Impression management	
Common followers	Impression management	Need for uniqueness
Common mutual followers	Trust, social bonding, impression management	Need for uniqueness

## 1.4 Model

Our objective is to evaluate the impact of network embeddedness on the propensity of a receiver to share content obtained from sender(s). We use a Cox proportional hazards model (Cox 1972) to estimate the hazard of sharing. In social networks, one challenge for a researcher is that a user may receive multiple feeds from different senders sharing the same content (or an aggregated feed from multiple senders) and the contribution of each co-sender on the decision to share is unclear.

At the consumer (receiver) level, a number of models have been proposed to deal with the impact of multiple senders (Toubia et al. 2014; Trusov et al. 2010) or multiple ad exposures (Braun and Moe 2013). A key difference between the present study and these studies is that our unit of analysis is a dyad rather than an individual. Individual level analysis often comes with some sort of aggregation on the sender side. For example, Aral et al. (2009) consider the overall effect of the number of shared friends on a user's likelihood to adopt a Facebook app, but the effect of individual friends' characteristics are not studied. Katona et al. (2011) accommodate multiple senders by considering the average characteristics of senders, which compromises model precision. While Trusov et al. (2010) do consider the effect of each individual sender on a user (restricted to be either 0 or 1), their model does not allow statistical inference on the effects of dyadic characteristics such as embeddedness. Sharara et al. (2011) focus on an adaptive diffusion model with the objective of establishing the effect of network dynamics on content sharing. They learn the "confidence values" between sender-receiver pairs based on past sharing for the purpose of making predictions. However, they do not deal with the estimation of the effect of dyadic characteristics on the propensity to share content.

Experimental studies (Aral and Walker 2012; Aral and Walker 2014) which conduct dyadic level analyses, avoid this problem by eliminating receivers getting notifications from multiple senders. While it eliminates the statistical challenge of dealing with multiple senders, it creates a controlled (and at times artificial) setting where the experiment inadvertently also controls for drivers of sharing that can be important in a natural setting of information sharing. For example, the need for

uniqueness is more likely to be a concern if multiple individuals in a user’s social network have shared the content as compared to a single individual sharing the content. We address this challenge by proposing a novel proportional hazards model that allows us to estimate the contribution of individual senders when multiple co-senders collectively cause a decision to share content.

#### 1.4.1 Dyadic Hazard

To ease model exposition, we present it in the context of sharing ad content over the social media platform, Digg.com (as it is the context of our primary dataset). On Digg, when a user (sender) diggs (shares) an ad (content), her followers (receivers) are immediately notified about her sharing activity in the form of a feed. A receiver can have multiple senders (co-senders) if more than one of her followees diggs the same ad. In addition to social feeds, users can also see the ad on the front page of Digg. Therefore, there are two types of shares on Digg: those driven by social sources (i.e., feeds from followees) and others driven by non-social sources (i.e., the front page). Other platforms such as Twitter have a similar process for information sharing between users connected in a social network.

Let  $i, j$ , and  $k$  index senders, receivers, and ads, respectively. Let  $t$  be the time elapsed since the creation of an ad. Let  $X_i(t)$  and  $X_j(t)$  represent the unitary attributes of sender  $i$  and receiver  $j$ , respectively (e.g., gender and activity level of a user on Digg). Let  $X_{ij}$  represent the dyadic attributes concerning sender  $i$  and receiver  $j$  (e.g., embeddedness measures),  $X_{ik}$  represent sender  $i$ ’s attributes that are specific to ad  $k$  (e.g., the time sender  $i$  diggs the ad  $k$ ), and  $X_{jk}$  represent receiver  $j$ ’s attributes that are specific to ad  $k$  (e.g., number of receiver  $j$ ’s followees that have shared ad  $k$ ). Let  $\lambda_{ijk}(t)$  represents the dyadic level hazard of sender  $i$  causing receiver  $j$  to adopt ad  $k$  at time  $t$ . Let  $\lambda_{k0}(t)$  represents the baseline hazard for ad  $k$ . The dyadic level hazard, stratified on ads, is given by

$$\lambda_{ijk}(t) = \lambda_{k0}(t) \exp\left(\beta_1 X_i(t) + \beta_2 X_j(t) + \beta_3 X_{ij}(t) + \beta_4 X_{ik}(t) + \beta_5 X_{jk}(t)\right), \quad (1.1)$$

$\lambda_{k0}(t)$  captures the baseline hazard for each ad. Note that the above semi-parametric formulation allows  $\lambda_{k0}(t)$  to change arbitrarily over time and across ads and allows us to capture static ad-specific effects such as the ad content and time-varying effects such as overall reduced tendency to share a specific ad with time. For example,  $\lambda_{k0}(t) = 0$  represents a case when an ad stops diffusing in the network. This formulation of dyadic hazard is similar to the formulations given in (Aral and Walker 2012; Aral and Walker 2014; Lu et al. 2013), but we allow one receiver to be exposed to the same ad from multiple senders.

Note that  $X_{ik}$  and  $X_{jk}$  include variables representing when a sender shares (which accounts for decaying effect) and the number of co-senders of a receiver, respectively. Due to users' need for uniqueness in online communities, we hypothesize that the effects of common followers and common mutual followers are negatively moderated by the popularity of information in H3 and H5. To test these effects, we consider interaction of the popularity of ads with common followers and common mutual followers and include these as dyadic attributes.

#### 1.4.2 Spontaneous Sharing

The basic specification of dyadic hazard ignores the possibility of users to share spontaneously. For example, a user may share content received from another user in the social network, or after receiving it directly from the platform or an external source. The latter type of sharing is termed as a spontaneous sharing and occurs via a non-social source (e.g., platform or external site). In order to incorporate the impact of non-social sources (e.g., the front page of Digg) in our study, we treat them as a special sender and use a dummy variable to capture their effect on the hazard rate:

$$\lambda_{ijk}(t) = \lambda_{k0}(t) \exp(\beta_0 s_i + \beta_1 X_i(t) + \beta_2 X_j(t) + \beta_3 X_{ij}(t) + \beta_4 X_{ik}(t) + \beta_5 X_{jk}(t)), \quad (1.2)$$

where the dummy variable  $s_i$  is 1 if the sender is the special sender and 0 otherwise. For the special sender, all undefined unitary and dyadic attributes are coded as missing and set to zero (or any other default value as the selection of default only affects parameter  $\beta_0$ ). The parameter  $\beta_0$  captures the combined effect of all non-social sources, as compared to a sender whose attributes

may be zero, on the sharing of the receiver. Since all users can adopt spontaneously, the special sender is a co-sender for every potential sharing user. Our dummy variable formulation enables us to seamlessly incorporate the effect of non-social sources.

### 1.4.3 Model Estimation

Let the parameter vector  $\theta = \{\beta_0, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5\}$  represent the entire set of parameters of our model. Let  $R_k(t)$  represent the set of receivers who have not shared ad  $k$  before time  $t$  (excluding), which is often referred to as the risk set. Let  $C_{jk}(t)$  represent the set of co-senders that have sent a feed regarding ad  $k$  to receiver  $j$  before time  $t$ . Let  $E$  represent the set of sharing events observed in the data and let  $E_{jk}$  represents the event of receiver  $j$  sharing ad  $k$ .

The key assumption of the proposed proportional hazard model is that the sharing of a receiver is collectively caused by all her co-senders, which is a standard assumption in previous non-dyadic models to deal with multiple senders (Toubia et al. 2014; Trusov et al. 2010) or multiple ad exposures (Braun and Moe 2013). In a hazard model, this means that the time it takes the receiver to share is determined by the overall hazard of the receiver. Assume that the hazards of the receiver to be influenced by each co-sender are independent conditional on the control variables, the overall hazard of receiver  $j$  to share ad  $k$  at time  $t$  is given by

$$\lambda_{jk}(t) = \sum_{i \in C_{jk}(t)} \lambda_{ijk}(t),$$

where  $\lambda_{ijk}(t)$  represents the dyadic level hazard of sender  $i$  causing receiver  $j$  to share ad  $k$  at time  $t$ . The additive form of the overall hazard results from the conditional independence assumption, which is a standard assumption for proportional hazards model.

Suppose event  $E_{jk}$  occurred at time  $\tau_{jk}$ , the partial log likelihood of this event can be written as

$$l(E_{jk} | \theta) = \ln P(E_{jk} | \theta) = \ln \left( \frac{\lambda_{jk}(\tau_{jk})}{\sum_{j' \in R_k(\tau_{jk})} \lambda_{j'k}(\tau_{jk})} \right) = \ln \left( \frac{\sum_{i \in C_{jk}(\tau_{jk})} \lambda_{ijk}(\tau_{jk})}{\sum_{j' \in R_k(\tau_{jk})} \sum_{i' \in C_{j'k}(\tau_{jk})} \lambda_{i'j'k}(\tau_{jk})} \right) \quad (1.3)$$

Note that the baseline hazard cancels out. The overall partial log likelihood of the entire dataset can then be written as

$$l(E|\theta) = \sum_{E_{jk} \in E} l(E_{jk}|\theta) = \sum_{E_{jk} \in E} \ln \left( \frac{\sum_{i \in C_{jk}(\tau_{jk})} \lambda_{ijk}(\tau_{jk})}{\sum_{j' \in R_k(\tau_{jk})} \sum_{i' \in C_{j'k}(\tau_{jk})} \lambda_{i'j'k}(\tau_{jk})} \right) \quad (1.4)$$

The parameters in our model can be estimated by maximizing the partial log likelihood given in Equation (1.4) using the Newton-Raphson method or other numerical optimization methods. In this paper, we use an enhanced Newton-Raphson algorithm to search for the optimal parameters of the partial log likelihood. Specifically, when the parameters reaches a non-concave region, we add a small positive number to the diagonal elements of the information matrix (typically slightly larger than the smallest eigenvalue of the information matrix in absolute value), as suggested by Schnabel and Eskow (1999), to make the information matrix positive definite. The effectiveness of the enhanced Newton-Raphson algorithm has been validated through extensive simulations. The above model collapses to the standard proportional hazards model when there is only one sender for each receiver.

Our proposed model has two advantages over prior specifications. First, it does not speculate on the contribution of each co-sender apriori, but allows the data to automatically determine the contribution of individual co-senders based on their characteristics. Second, it is applicable even if only some of the co-senders have a significant impact on the sharing, as the likelihood in Equation (1.3) essentially captures the probability of the true cause belonging to the set of co-senders. Lacking information on which subset of co-senders have real effects will increase the standard errors of the parameter estimates, but will not bias the point estimates. In Appendix 1.1, we show using simulations that the proposed model can recover the true parameters with negligible errors, regardless of whether the sharing events are caused by all co-senders collectively or only one of the senders. In contrast, we find that models that make assumptions on the contributions of co-senders apriori can result in substantial bias (see Table A1.1 in Appendix 1.1).

#### 1.4.4 Identification

A primary challenge for determining the impact of the network characteristics on user actions is that the results could be biased due to unobservable characteristics. For example, a sender with high popularity offline might be more influential than other senders with similar online characteristics. While such offline information might be observable to the receiver, it is often unknown to the researcher. Similarly, a receiver with stronger interest in ad-related content might be more likely to share ads in general, and such topical interest of individual receivers is often not available to the researcher. Missing information on either senders or receivers can bias model estimates. To address this concern, we allow for random effects at the sender-level<sup>4</sup> and the receiver-level, which allow each sender and receiver to have a random intercept that captures the main effect of unobserved characteristics. We also consider random effects at the dyadic level to account for dyad-specific unobservables, following previous studies in network contexts (Hoff 2005; Lu et al. 2013; Narayan and Yang 2007). Note that it is possible that the unobserved characteristics are correlated with observed characteristics. For example, a sender with high unobservable popularity may also have lot of connections and, as a result, a larger overlap with the receiver's connections as compared to a less popular sender. As random effects cannot accommodate such correlations, we estimate models with fixed effects at the sender level (fixed effects allow for unobserved characteristics to be correlated with observed characteristics).

In addition to unobserved characteristics, two additional concerns for identification are spontaneous shares and endogenous communication patterns (Aral and Walker 2014). For the former, we explicitly control for the possibility of spontaneous shares, by treating all non-social sources as a special sender. Such a control not only teases out the effect of non-social sources, but also alleviates, to some extent, the concern that a receiver is sharing due to her inherent propensity to share. For the latter, in our application, the platform sends a notification to all followers of a sender.

---

<sup>4</sup> Given that the special sender representing the effects of non-social sources is intrinsically different from other senders, we allow the variance of the frailty term for the special sender to be different from other senders.

Thus, there is no selection bias on who can see the content (i.e., no endogenous communication patterns).

A fourth problem with identifying information sharing across a dyad is that a receiver often sees the same information from multiple senders before sharing, and the quantitative contribution of each co-sender may be unclear. We address this challenge statistically by proposing a novel proportional hazards model that determines the contribution of each co-sender based on their characteristics.

## 1.5 Data

We seek to understand how embeddedness between a sender and a receiver connected in a social network impacts the sharing behavior of the receiver. A dyadic level study imposes stringent requirements on the data. First, we need a sample of marketing-related messages or content generated on a social media platform by firms.<sup>5</sup> Next, for each piece of content, we need complete information regarding how the content propagates through the network from activated users (senders) to their followers (receivers). Such information includes the profile and social graph information of all involved users (both senders and receivers), as well as time-stamped sharing information at the individual user level. The sample of involved users can be identified by traversing the audience of activated users progressively. Specifically, we can start from a set of seeds (e.g., the author or users who spontaneously share content) and then treat the followers of these seeds as receivers. This process iterates when a receiver become activated, i.e., she shares the content, until the end of the observation time window. This progressive user sampling approach based on ego's network allows us to focus on users who are relevant to our analysis. A similar approach has been employed by other researchers interested in the effects of dyadic network characteristics (Aral and Walker, 2014; Shi et al. 2014). The set of users chosen by the progressive sampling approach are all the activated users (senders) and their followers (receivers). Finally, the profile and social graph information on these users can be collected retrospectively from historical data on social

---

<sup>5</sup> This is important as we can establish the implications of our results for firms utilizing social media to reach out to consumers.

media platforms. Note that if there are users with regular exposures to non-social sources (e.g., portal pages), we can also consider them as receivers.

We collected a dataset with the desired information from Digg.com, one of the largest online social news aggregation websites. On the website, users can highlight (“digg”) their favorite content and the activity is visible to all of their followers. Digg introduced a native advertising model, called diggable ads, in 2009, which remained on the website until Digg’s acquisition in August 2012. The feature allowed an advertiser to promote sponsored content in the feeds of Digg’s users. Diggable ads were seamlessly integrated with organic stories and displayed at three fixed positions of the eighteen slots available on the front page. At the time we collected the data, Digg maintained a Twitter like social network structure (see footnote 1), allowing users to follow each other.

Initially ads are only shown on the front page. Users can digg up or down an ad after viewing it just like digging an organic story. In that case, the ad is also included in the news feed of all their followers including mutual followers. Other users can explore the ad on the front page or navigate through feeds of their followees’ activities in the “My News” page. All activities associated with an ad are automatically combined into a single feed for clarity. The identities of the involved followees are displayed side by side in the combined feed. Due to this feed combining feature, it is likely that each followee (co-sender) more or less has some effect on the activity of the focal user (receiver). Diggable ads were identical to organic stories except for an inconspicuous flag “sponsored by xx” below them. Diggable ads are removed from the front page when the associated advertiser runs out of budget, but users can still see them from social feeds.

We investigate the sharing of diggable ads.<sup>6</sup> For the purpose of this study, we focus on all ads (31) created during a randomly chosen month-long period (May 24th, 2012 to June 25th, 2012). As

---

<sup>6</sup> Identification of the effect of network drivers is easier for diggable ads as opposed to that for organic content. Diggable ads are guaranteed to be displayed on the front page before running out of budget, whereas whether an organic story is displayed on the front page depends on many factors, including the diggs the story receives and the freshness of the story. Therefore, the spontaneous hazard of organic stories may change radically over time due to their unstable visibility on the front page, which makes it difficult to mod

mentioned earlier, we need the profile and social graph information of all involved users in the ad sharing process to study the effect of embeddedness associated with dyads on the sharing behavior. In the Digg setting, since all users can see the ads from the front page, they are all potential receivers. In order to control the size of our dataset, we only consider active users who can potentially digg or share these 31 ads.<sup>7</sup> We define a user as active if she has dugg at least one ad in the past and still maintained some activity on Digg such as posting, digging and commenting other content in the focal time period.<sup>8</sup> In robustness analysis, we also consider users who have dugg an ad in the past but have no activity during the focal period and find that our results are similar.

For each potential receiver, we generate one dyadic observation for her if one of her followees shares the ad. Since everyone has access to the front page<sup>9</sup>, we generate one additional dyadic observation for each potential receiver, with the front page being the sender. The act of digging allows the user to share the ad with her followers. One converts from a receiver to a sender immediately after the sharing activity, implying that senders are a subset of receivers. A nuanced issue in our context is that ads stop showing up on the front page after a certain period and as a result, the spontaneous hazard becomes zero. To ensure that is not the case, we choose a censoring time for each ad as the last time when the ad was shared spontaneously by a user. The censoring time of an ad ranges from 1.4 days to 7 days, after its creation. The average censoring time is approximately 5 days. This resulted in a sample of 8,164 users and 95,144 dyads. Table 1.4 shows the summary information of the dataset. The table shows that 32% of shares have more than one co-sender (excluding the special sender “front page”), and the average number of co-senders is 2.82, including the front page.

---

<sup>7</sup> Focusing on active users allows us to remove inactive users who are not at risk of sharing anymore. In practice, marketers often focus on such high risk users in their targeting campaigns (e.g., sending coupons to customers who have purchased their products in the past or who have met some threshold on the amount spent).

<sup>8</sup> We have access to profile information of all users who ever dugg one of the diggable ads between October 2010 and July 2012, including gender, location, number of diggs, number of comments, number of submissions, number of followers, and number of followees.

<sup>9</sup> On Digg, the front page is the primary non-social source for the sharing of ads. Another possibility, albeit rare in our context, is that users may discover the content through a search engine. For ease of exposition, we refer all non-social sources as the front page.

Table 1.4 Summary Statistics

Number of ads/tweets	31
Number of sharing user (senders)	1,058
Number of potential sharing user (receivers)	8,164
Number of <sender, receiver> dyads	95,144
Number of <sender, receiver, ad> tuples	560,044
Number of spontaneous tuples	222,846 (40%)
Number of social tuples	337,198 (60%)
Number of shares (diggs)	2,810
Number of spontaneous shares	1,438 (51.2%)
Number of potential influenced shares	1,372 (48.8%)
Percentage with more than one co-senders (excluding special sender)	32.1%

We used the APIs provided by Digg to collect the social graph of all potential users who could share the sample ads. Due to the rate limit on API calls, it took 19 days (June/7/2012- June/26/2012) to collect a single snapshot of the complete set of followers and followees for these users. One concern with this data extraction process is that network of users may have changed even during the sample period. However, the extent of network changes is small in our setup. By comparing the profile information of users on June 7 and July 9, we found that the both the follower and followee numbers changed less than 5% on the log scale for 85% of users and the mean absolute relative change on the log scale is less than 2.5%. Thus, changing network is not likely to significantly impact our results. As a further check, we split our sample in two subsamples and repeat the analysis for each (see robustness checks). All our substantive findings are robust.

We use several control variables pertaining to the sender, the receiver, and the sender-receiver dyad. These variables, summarized in Table 1.5, include the unitary network attributes of the sender/receiver, the engagement level of the sender/receiver, the demographics of the sender/receiver, the timing of the sender's share, the number of co-senders in the receiver's network, and so forth. Table 1.6 summarizes the summary statistics for the main unitary and dyadic network attributes and control variables.

Table 1.5 Descriptions of Independent Variables

Independent Variable	Description	
$X_i/X_j$	<b>Attributes of sender <math>i</math> / receiver <math>j</math></b>	
Network attributes	followees	Number of followees (out-degree)
	followers	Number of followers (in-degree)
	mutual	Number of mutual followers
Engagement levels	diggs	Total number of diggs
	comments	Total number of comments
	submissions	Total number of submissions
	avgDiggs	Average number of diggs per month
	avgComments	Average number of comments per month
	avgSubmissions	Average number of submissions per month
Others	gender	Male, female, or missing
	isSocial ( $s_i$ )	1 if sender $i$ is a social source (i.e., followee), otherwise 0
	isSubmitter	1 if the sender is the submitter of the ad, otherwise 0
$X_{ij}$	<b>Attributes of a sender-receiver dyad</b>	
Dyadic network attributes	isMutual	Does the sender and the receiver follow each other mutually
	commonFollowees	Number of followees shared by the sender and the receiver
	commonFollowers	Number of followers shared by the sender and the receiver
	commonMutuals	Number of mutual followers shared by the sender and the receiver
$X_{ik}$	<b>Sender-specific attributes of an ad</b>	
Sharing timing	wday	Day of a week when sender $i$ dugg ad $k$
	hour	Hour of a day when sender $i$ dugg ad $k$
	shareTime	Hours taken for sender $i$ to adopt since the creation of ad $k$ , 0 for the front page
$X_{jk}$	<b>Receiver-specific attributes of an ad</b>	
	co-senders	Number of followees (co-senders) of the receiver who have already shared
$X_k$	<b>Attributes of ads <math>k</math> (only interaction with other variables can be identified)</b>	
	popularity	Number of diggs on an ad at a given time point

Table 1.6 Key Statistics of Main Variables

	Zeros	Mean	SD	Min	Median	Max
<b>Unitary Network Attributes of All Users</b>						
Number of followees	141	268.0	423.7	0	118	10122
Number of followers	146	386.3	1091.0	0	136	29331
Number of mutual	424	114.4	203.8	0	36	4598
<b>Dyadic Network Attributes of Sender-Receiver Dyads</b>						
isMutual (1 – reciprocal, 0 – non-reciprocal)	63733	0.27	0.44	0	0	1
Number of common followees	4736	41.1	52.5	0	23	814
Number of common followers	2182	100.7	334.2	0	26	9812
Number of common mutual followers	19805	17.1	35.3	0	4	594
<b>Popularity of Ads</b>						
Number of diggs	0	93.4	86.2	4	95	295

Table 1.7 outlines the correlation among dyadic network characteristics. As discussed earlier, to clearly identify the effects of different overlapping connections, we exclude common mutual followers when counting the number of common followees and common followers. The correlations among the three embeddedness metrics are not very high and suggest that these metrics are capturing different drivers. Further, the estimates of the correlated variables were stable with changes in model specifications and data samples, suggesting that multicollinearity is unlikely to be an issue.

Table 1.7 Correlation among Dyadic Network Characteristics

	isMutual	logCommonFollowees	logCommonFollowers	logCommonMutuals
isMutual	1.00	0.16	0.07	0.53
logCommonFollowees		1.00	0.53	0.46
logCommonFollowers			1.00	0.46
logCommonMutuals				1.00

In order to understand how ads were shared over time, we plot the Kaplan-Meier survival curve for some sample content (see Figure A1.1 in Appendix 1.3). Note that the sharing activities on most ads basically ceased at the censoring time. The sharing graphs for two sample ads with average popularity are shown in Figure A1.2 in Appendix 1.3. These graphs demonstrate that path length is short (around 2 on average) for content as they propagate through the user network. This is in agreement with the observation made by Goel et al. (2012) about short path lengths for diffusion in online social networks. Note that our model assumes that the effect of co-senders can either increase or decrease. This may not accurately capture the aggregate diffusion pattern especially when the network is saturated and the effect of co-senders is very likely to decrease. However, path lengths for our data suggest that the network is not saturated and alleviates such concern. Next, we discuss our results on the role of embeddedness on the sharing propensity of the receiver.

## 1.6 Results

### 1.6.1 Main Results

Table 1.8 summarizes the results of four model specifications.<sup>10</sup> Our main model of interest is model 4 that includes interaction terms representing the moderating effect of ad popularity on common followers and common mutual followers. We have also estimated models with no interaction terms or including only one of the two interaction terms (models 1-3, respectively). Likelihood ratio tests suggest that model 4 is preferred over models 2 and 3 ( $p < 0.05$ ). The following discussion is based on the estimates from model 4 unless otherwise specified.

Table 1.8 Parameters Estimates of Different Model Specifications

	Model1	Model2	Model3	Model4
<b>Embeddedness</b>				
logCommonFollowees	0.23***	0.174**	0.175**	0.175**
logCommonFollowers	0.845***	1.364***	0.829***	1.074***
logCommonMutuals	-0.245***	-0.19***	0.799***	0.418**
<b>Interactions with Popularity</b>				
logCommonFollowers:logPopularity		-0.153***		-0.071**
logCommonMutuals:logPopularity			-0.258***	-0.16***
<b>Fitness</b>				
logLikelihood	-22661	-22623	-22620	-22618
AIC	45401	45326	45320	45317

Significance levels:  $p < 0.001$  (\*\*\*),  $p < 0.01$  (\*\*),  $p < 0.05$  (\*), and  $p < 0.1$  (.). The main effect of logPopularity cannot be identified as everyone sees the same digg number at a given time point, the effect of which is cancelled out in the likelihood. Model 2 is chosen as our main model based on fitness.

*Common followees.* The number of common followees has a positive effect on the sharing propensity of the receiver. This validates H1. The number of common followees reflects the similarity between the sender and the receiver's tastes and expertise. For the purpose of impression management, users tend to share content representing their taste or expertise (Berger 2014; Packard & Wooten, 2013). Thus, the more common followees the receiver has with the sender, the more likely the receiver will also share the content from the sender. Note that we obtain

<sup>10</sup> We omit the coefficients on control variables for clarity. Please see Appendix B for the complete set of parameter estimates.

this result after controlling for the effect of common mutual followers, which represent close friends. Thus, our result suggests that common followees can also be used to capture similarity or homophily between users (McPherson et al. 2001).

*Common followers.* The simple effect of common followers (when the logarithm of the content popularity is zero) is positive, suggesting that the number of common followers has a positive effect on dyadic influence when the popularity of ads is low. This finding validates H2. As discussed earlier, the number of common followers reflects the similarity between the sender and receiver's audiences. Users tend to share content of interest to their audience to impress them (Aral and Van Alstyne 2011; Wu et al. 2004). Therefore, if the receiver has a similar audience with the sender, the receiver is likely to make the same decision as the sender (i.e., to share), especially when the content is relatively novel and the concern around uniqueness is not strong. The negative interaction of common followers with content popularity confirms H3: the effect of common followers decreases with content popularity, validating users' need for uniqueness in content sharing (Ho and Dempsey 2010). This is similar to extant findings that indicate that consumers with a high need for uniqueness may decrease the consumption of a product if it becomes commonplace, also known as the reverse-bandwagon effect (Cheema and Kaikati 2010; Granovetter and Soong 1986).

*Common mutual followers.* The simple effect of common mutual followers (when the logarithm of content popularity is zero) is positive and demonstrates that, when the content is relatively novel, common mutual followers has a positive impact on sharing. This finding validates H4. The existence of neighbors mutually connected to two users expands the bandwidth of communication between them and increases their trust in each other (Aral and Van Alstyne 2011; Burt 2001). The negative interaction of common mutual followers with popularity confirms H5. This finding shows a boundary condition for the positive effect of embeddedness previously reported in undirected networks (Aral and Walker 2014; Bapna et al. 2015). Specifically, the effect of common friends might be positive only when the information to be communicated is relatively novel (or not as popular).

In sum, all our proposed hypotheses find support from data. Our results show that the effect of embeddedness in directed networks varies across different types of “neighbors”. Moreover, the impact of common followers and common mutual followers are negatively moderated by content novelty. The interaction effects suggest that users are eventually going to cease sharing due to concerns around uniqueness. As a result, the content is likely to diffuse for short distances within a network. This may explain the short information cascades reported in literature (Goel et al. 2012) and also observed in our dataset (Figures W1 and W2).

In addition to the findings on the three embeddedness metrics, it is worthwhile highlighting the estimates on two additional variables (i.e., co-senders and shareTime), which help us understand how each co-sender contributes to a receiver’s propensity to share. First, the effect of co-senders is negative, showing that the marginal effect of a co-sender decreases with the number of co-senders (though the overall effect of all senders may increase). Second, the effect of shareTime is positive<sup>11</sup>, suggesting that the later a co-sender shared, the stronger effect the co-sender has on the receiver. This documents a recency effect for co-senders consistent with previous findings that social effects decay over time (Bakshy et al. 2012; Haenlein 2013; Nitzan and Libai 2011; Trusov et al. 2009).

### 1.6.2 Robustness Checks

*Unobserved Characteristics.* A potential concern with our analysis is that sharing of content could be driven by unobserved characteristics at the sender, the receiver, and even the dyad level. The dyadic observations with the same sender, receiver or dyad may not be independent because of common unobserved characteristics. As a robustness check, we consider sender-specific, receiver-specific and dyad-specific random effects. We also account for the effects of unobserved characteristics with a fixed effects approach as it allows for unobserved characteristics to be

---

<sup>11</sup> It can be easily proved that, in a proportional hazards model, using shareTime (i.e., how long did it take for a sender to adopt) is equivalent to using recency (i.e., how long ago did the sender adopt), because the sum of the two variables equals the time elapsed since the creation of the ad. The only difference is that the estimates on both variables will have opposite signs. We use shareTime as it does not vary over time, which facilitates the estimation.

correlated with observed characteristics. While the fixed effects approach appears to be more flexible than the random effects model in terms of its assumptions, it is more sensitive to the issue of insufficient reoccurrence. Specifically, in the proportional hazards modeling framework, a random effects approach tends to provide more reliable estimates than the fixed effects approach as the former penalizes large individual effects (Therneau 2000) and prevents the model from over-fitting. With that being said, we still estimate fixed effects on the sender level but not on the receiver-level as the low reoccurrence frequency of receivers in our data may result in substantial incidental parameter bias in the estimates (Allison 2002; Lancaster 2000). Fixed effects on the dyadic level are not a viable alternative as well, as then the effects of dyadic network characteristics are not identified. Note that the random/fixed effects allow us to account for unobserved factors such as the fact that some users might be bots on Digg.

Table 1.9 presents the results from different models with random and fixed effects at sender, receiver and dyad levels. Overall, the estimates on the dyadic network characteristics are qualitatively similar across different model specifications.

Table 1.9 Parameters Estimates from Different Random/Fixed/Mixed Effects Models

	none	rs	fs	rs-rr	fs-rr	rs-rr-rd	fs-rr-rd
<b>Embeddedness</b>							
logCommonFollowees	0.175**	0.146.	0.233***	0.2*	0.165***	0.118	0.152*
logCommonFollowers	1.074***	1.077***	1.078***	0.674***	0.658***	1.784***	0.861***
logCommonMutuals	0.418**	0.364.	0.193	0.679**	0.361*	1.023***	1.119***
<b>Interactions with Popularity</b>							
logCommonFollowers:logPopularity	-0.071**	-0.068*	-0.07**	-0.054.	-0.082**	-0.166***	-0.34***
logCommonMutuals:logPopularity	-0.16***	-0.153***	-0.132***	-0.178***	-0.098*	-0.226**	-0.186**
<b>Fitness</b>							
logLikelihood	-22618	-22538	-22278	-19974	-19628	-20226	-19960
AIC	45317	45163	46727	40038	40932	40544	41597

In row 1, the first letter represents whether fixed (f) or random (r) effects is used. The second letter indicates the subject ("s" for sender, "r" for receiver, and "d" for dyad) on which the specified effect is applied. Therefore, "rs" represents a model with random effects on sender, and "fs-rr-rd" represents a model with fixed effects on sender, random effects on receiver, and random effects on dyad. "rs" is the main model used in this paper. The model "none" doesn't include random or fixed effects on any subject.

*The Growth of Network Structure.* Another concern with our analysis is that the network structure among users may change over time but we used a static snapshot. Note that Digg users often establish new ties but rarely break old ties. The direct consequence of the inaccurate network structure information is that the number of observed co-senders for a receiver could be larger or smaller than the actual number of co-senders for the receiver, depending on whether the receiver dug the ad before or after the time her network information was collected by us. In our dataset, almost all the ads were posted on three days: May 24, June 1, and June 25. In order to test the sensitivity of our results to this issue, we split the dataset into two subsets: one focusing on ads created between May 24 and June 1, and another focusing on ads created on June 25. Recalling that the network structure is collected during June 7- June 26, the number of co-senders is likely to be overestimated on the first dataset as the network structure is collected afterwards. In the second dataset, the number of co-senders is likely to be underestimated as most of the digging activities take place after the network structure is collected. If overestimation or underestimation of the number of co-senders causes a substantial bias on our estimates, the results on these two subsets should be very different from that on the full dataset. Table 1.10 summarizes the results on the two subsets, respectively. The results show that the estimates on the two subsets are highly consistent with that on the full dataset.

Table 1.10 Parameters Estimates on the Two Subsets

	May24-June1 (1879 Events)	June 25 (931 Events)
<b>Embeddedness</b>		
logCommonFollowees	0.233***	0.204*
logCommonFollowers	0.858***	1.147***
logCommonMutuals	0.654***	0.643**
<b>Interactions with Popularity</b>		
logCommonFollowers:logPopularity	-0.055.	-0.038
logCommonMutuals:logPopularity	-0.186***	-0.274***
<b>Fitness</b>		
logLikelihood	-15149	-7427
AIC	30379	14935

*Inactive Users.* In our main analysis, we only consider active users as candidates for sharing. We also re-estimate our model by including data for users who have dugg an ad in the past but are not active during the panel period. Results are included in Table A1.3 in Appendix 1.4 and are qualitatively similar to our main analysis.

### 1.6.3 Generalizability to Other Social Networks

To test whether our findings generalize to other directed networks, we collected an additional dataset from Twitter. In the context of Twitter, the act of sharing is retweeting. Similar to Digg, the sharing is spontaneous if a user shares a tweet before any of her followees do. Otherwise, the sharing is considered as sharing influenced by others. To make sure that the content of the Twitter dataset is similar to that of the Digg dataset and also to improve the managerial relevance of our study, we focus on the sharing of brand-authored tweets.

We focus on nine brands listed by Fortune magazine as the top fortune 500 companies using social media.<sup>12</sup> We first collect the tweets authored (or retweeted in rare cases) by these brands in the past 10 days.<sup>13</sup> Then for each tweet, we collect the social graph information needed for our analysis retrospectively in two steps. As the first step, we collected the social graph information of all retweeters (including the author) of the tweet. These users represent the set of senders for the focal tweet. Next, we collected the social graph information for the followers (receivers) of the senders. Since the density and network size of Twitter users is much higher than that of Digg users<sup>14</sup>, collecting data for all followers of every sender is not feasible due to API restrictions.<sup>15</sup> In order to control the data size, for every sender, we consider all followers who retweet. However, we randomly sample other followers from the sender's ego network using the risk set sampling approach (Langholz and BORGAN 1995; Langholz and Goldstein 1996). Specifically, depending on popularity of each brand, we sample 5~20 followers from the ego network of each sender

---

<sup>12</sup> <http://fortune.com/2014/06/02/500-social-media/>

<sup>13</sup> We collected two sets of tweets for each brand in about six weeks.

<sup>14</sup> In our sample, a user on Digg, on an average, has around 400 followers whereas a user on Twitter has around 19000 followers.

<sup>15</sup> <https://dev.twitter.com/rest/public/rate-limits>

(sample size is smaller for popular brands with more data to collect).<sup>16</sup> We then collect the profile information for all the identified users. Similar to the Digg dataset, we focus on the receivers who are still active in the past three months. We focus on 4740 sharing activities on 74 tweets with more than 20 retweets in our analysis.<sup>17</sup> Further description and statistics for the Twitter dataset are shown in Tables A1.4-1.7 in Appendix 1.4. Table W6 shows the complete set of results for the Twitter dataset.

Table 1.11 Parameter Estimates on Twitter Dataset

	Model1	Model2	Model3	Model4
<b>Embeddedness</b>				
logCommonFollowees	0.294***	0.311***	0.299***	0.309***
logCommonFollowers	-0.127***	0.227***	-0.132***	0.144***
logCommonMutuals	-0.035	0.013	0.643***	0.284***
<b>Interactions with Popularity</b>				
logCommonFollowers:logPopularity		-0.105***		-0.081***
logCommonMutuals:logPopularity			-0.174***	-0.075***
<b>Fitness</b>				
logLikelihood	-29378	-29324	-29340	-29319
AIC	58822	58717	58747	58708

Table 1.11 summarizes the parameter estimates for the three embeddedness metrics for Twitter dataset. Our main model of interest is model 4 and has the best fit. The results show that the findings on the Twitter dataset are consistent with that on the Digg dataset. The coefficient of common followees is positive and significant. The coefficients of common followers and common mutual followers are also positive and significant. And, the coefficients of the terms capturing interaction of these variables with popularity are negative and significant. This pattern of results demonstrates the generalizability of our findings from Digg to other directed social media platforms like Twitter. Unlike in the Digg dataset, however, we cannot effectively estimate random/fixed

<sup>16</sup> In order to ensure that the number of followers sampled does not affect our results, we tried to increase the sample size to as many as 50 followers for each retweeter on some brands and find the estimates are rather robust to the sample size.

<sup>17</sup> We also tried using tweet samples with other popularity thresholds, such as 5, 10, 30 and 40 and our results are qualitatively similar.

effects on the Twitter dataset as the reoccurrences of each sender, receiver, and dyad are substantially lower.

It is important to highlight that there are a few differences in how we collect and analyze the Digg and Twitter datasets, mainly to incorporate the contextual differences between the two platforms. The first difference is that, in the Digg dataset, we treat all users as candidates for spontaneous sharing of an ad, as they all can see the ad on the front page of Digg. In the Twitter dataset, however, for each tweet, only the followers of the author (i.e., the brand) or retweeters are candidates for spontaneous sharing because there are no such non-social sources like front page that guaranteed substantial exposure for non-followers. Second, in contrast to Digg, Twitter often only shows the feed from the earliest co-sender to the receiver and does not provide any clue about the other co-senders' activity on the same tweet. However, our model can effectively handle the case when only one of the co-senders has a significant impact. Therefore, this should not bias our estimates, especially given that only 7% of retweeters in our sample have more than one co-sender. What is noteworthy is that despite these differences between Digg and Twitter, we obtain highly similar results and it further demonstrates the generalizability of our findings.

## **1.7 Discussion & Conclusion**

Social media platforms hold the potential to reshape the manner in which consumers generate, spread and consume content. Understanding what leads to effective information sharing at the dyadic level lies at the core of cost-effective content propagation on these platforms. While the effects of unitary network attributes have been well-studied in the literature, studies on the effects of dyadic network attributes on information sharing are nascent and predominantly focus on undirected networks.

In this paper we study the effect of a dyad's network embeddedness on information sharing in directed networks. More specifically, we quantify the effects of common followees, common followers, and common mutual followers between a sender and a receiver on the propensity of

sharing by the receiver. Substantively, our results show that the effect of embeddedness in directed networks varies across different types of “neighbors”. The number of common followees is positively associated with receiver’s propensity of sharing. Other embeddedness measures such as number of common followers and common mutual followers also have positive effect on this propensity. However, the latter positive effect decreases with the popularity of shared content. Thus, our study provides insight into consumer behavior in online information sharing and adds to the existing literature highlighting the role of uniqueness in social consumption (Cheema and Kaikati 2010; Zeng and Wei 2013). It is possible that uniqueness concerns may be preventing users from sharing the information received from others once the information becomes less novel. This in turn might be causing small cascades. Thus, our results provide a potential explanation for the relatively small size of information cascades that have been observed in online social networks (Goel et al. 2012)

We make a methodological contribution as well by proposing a new hazard rate modeling approach to more accurately determine the contribution of individual senders on influencing a receiver when multiple senders are involved. Quite often, consumers may respond only after the content is seeded by multiple senders (Centola and Macy, 2007). Even if detailed tracking information is available for each user, it would be difficult to determine the exact contribution of each sender in the content sharing process.<sup>18</sup> Previous work either makes strong assumptions about how the contribution should be attributed to different senders (Aral et al. 2009; Braun and Moe 2013; Katona et al. 2011; Toubia et al. 2014; Trusov et al. 2010) or does not focus on the identification of the effect of shared characteristics (Sharara et al. 2011; Trusov et al. 2010). Our approach makes no such assumptions and, as a consequence, can help to better tease apart the effect of the shared network attributes.

For marketing managers, we provide insights on how to target customers in a directed network at a micro level. Many platforms support micro level targeting to improve the efficacy of targeting (e.g.,

---

<sup>18</sup> While a platform can track the actual time when a receiver sees content from one or more senders and the sequence in which the content is received, it cannot determine how consumer is weighing these different feeds in her decision to adopt the content and in turn send it to her followers.

display of promoted tweets on Twitter) and prevent information overload for their members (e.g., filtering of feeds on Weibo). Our results show that platforms such as Twitter or Weibo can improve their targeting or filtering by focusing on dyads embedded in different types of connections (i.e., followees, followers, mutual followers). As a concrete example, when deciding whether or not to show a promoted tweet to a given user<sup>19</sup>, Twitter may want to consider how many common neighbors this user shares with the author, as well as the overall popularity of the tweet. Specifically, targeting users who have more common followees with the author can be more effective. Targeting users who have large numbers of common followers and common mutual followers can also be effective when the tweet is not that popular, but might be counterproductive when the tweet is already sufficiently popular. Finally, as compared to most previous studies that primarily focus on the sharing of organic content in social networks, the analysis of this paper is based on the sharing of sponsored ads and brand-authored tweets, which makes our findings of direct relevance to marketers.

Our work can be extended in several ways. First, it is likely that characteristics of the content can influence how much it is shared within dyads (Berger and Milkman 2012). Our modeling framework allows us to account for the heterogeneity of content but it would be useful to understand if the magnitude or direction of our results is sensitive to type of content being shared. Further, we considered sponsored ads and brand-authored tweets. It is possible that the user behavior may be different for organic content. Future studies should investigate the role of content characteristics in moderating the effect of network attributes on information sharing. Second, from a modeling standpoint, we did not have information on whether or not a user actually saw the feed. Without the impression information, we are essentially modeling the overall hazard of a user to read and adopt an ad. This coarse modeling structure may increase the standard errors of our estimates. However, the impression information is typically only known to social media platforms. Future research should

---

<sup>19</sup> Once a tweet is promoted, Twitter can display the tweet to any user on the platform, even though this user doesn't follow the author of the tweet. However, in practice, to avoid spamming users, Twitter only displays promoted tweets to selective users deemed relevant. Note that an advertiser can promote a tweet authored by a random user.

explore alternative approaches to address the lack of impressions such as conducting experiments where such information can be obtained from users (De Bruyn and Lilien 2008) or developing a latent model to capture the effect of impressions (Kang et al. 2013). Finally, the assumption that the existence of one co-sender does not cannibalize or reinforce the effects of other co-senders is restrictive. In our analysis, we address this problem by allowing the hazard of a co-sender to change with the number of co-senders (i.e., shared followees in Table A1.2 of Appendix 1.2). The negative coefficient on shared followees suggests that the marginal effect of a co-sender decreases with the number of co-senders (i.e., the cannibalization effect exists). However, this remedy strategy may not be satisfactory if the hazards of individual co-senders change by different multiplicative scales as the number of co-senders increase. Future studies should explore the non-linear effect of the number of co-senders on the outcome.

## 2. Participation vs. Effectiveness of Paid Endorsers in Social Advertising Campaigns: A Field Experiment

### 2.1 Introduction

Social advertising leverages social connections among consumers to reach and influence a target audience. This business practice is becoming increasingly popular. According to BI intelligence<sup>20</sup>, social advertising spending in the US will top \$8.5 billion in 2015 and reach nearly \$14 billion by 2018. Globally, it is expected to reach \$23.7 billion in 2015 and \$36 billion by 2017, capturing 16% share of all digital ad spending.<sup>21</sup> Two thirds of marketers believe that social media is core to their business, and 70% of them plan to increase the budget on social media marketing.<sup>22</sup>

The prevalent social advertising mechanism is a centralized system in which advertisers submit ads to social media platforms (publishers) who then decide how to distribute the ads. Two drawbacks of this system are that advertisers have no direct control over the selection of endorsers (e.g., users who share/retweet an ad on Facebook/Twitter) and that endorsers are not incentivized to get engaged. Paid endorsement, in contrast, is a decentralized mechanism that allows advertisers to bypass publishers and recruit individual endorsers of their own choice at pre-specified prices. Specifically, advertisers post tasks on a paid endorsement platform (a broker website similar to Amazon Mechanical Turk) and microbloggers registered on the platform can take on the tasks requiring them to post or retweet some ad for monetary rewards. Paid endorsement has gained particular popularity in China, with many websites acting as platforms for paid endorsement. Weibo.com, the largest Chinese microblog site with more than 500 million users, launched its official paid endorsement platform in 2012.

---

<sup>20</sup> <http://www.businessinsider.com/social-media-advertising-spending-growth-2014-9>

<sup>21</sup> <http://www.emarketer.com/Article/Social-Network-Ad-Spending-Hit-2368-Billion-Worldwide-2015/1012357>

<sup>22</sup> <http://www.adweek.com/socialtimes/social-marketing-2015/504357>

Despite the growing interest in paid endorsement and social advertising in general, its effectiveness remains in question. Two thirds of advertisers are uncertain about the effectiveness of social advertising.<sup>23&24</sup> The effectiveness of a paid endorsement campaign depends on how many endorsers participate and on how well they expand reach (i.e., views), generate engagement (i.e., likes, comments, and retweets), increase traffic (i.e., clicks), and boost sales.

A key question facing marketers is how to incent endorsers. One problem with paid endorsement is that the incentive of participants is contingent on participation rather than performance, as monitoring performance is often practically infeasible or too costly. So far, the incentive on most paid endorsement platforms (e.g., weizuitui.com and sandaha.com) are simply determined by the number of followers an endorser has. Research on survey response behavior shows that incentives unrelated to performance typically increase participation, but rarely affect performance (Cantor et al. 2008; Singer and Ye 2013). The same may hold in paid endorsement campaigns. A second source of complexity is that the reaction to a particular level of financial incentive is likely to vary across potential endorsers with different award histories, as implied by prospect theory positing reference dependence and loss aversion, two phenomena well-documented to affect consumer behavior (Greenleaf 1995; Hardie et al. 1993; Kalyanaram and Winer 1995; Lattin and Bucklin 1989).

A second key question is which endorsers to target. Whether or not an endorser is worth targeting not only depends on the endorser's effectiveness in generating desired outcomes (e.g., engagements and sales), but also the endorsers' willingness to participate, as only participants can generate real outcomes. To design successful targeting strategies, it's critical for marketers to understand which endorsers are responsive (in participation) and which endorsers are effective (in generating outcomes), and more importantly, whether responsive endorsers are also effective. Meanwhile, given that different marketers may have different objectives in their campaigns, whether

---

<sup>23</sup> <http://www.nielsen.com/content/dam/corporate/us/en/reports-downloads/2013%20Reports/Nielsen-Paid-Social-Media-Adv-Report-2013.pdf>

<sup>24</sup> <http://www.socialmediaexaminer.com/SocialMediaMarketingIndustryReport2014.pdf>

the effectiveness of endorsers varies across different types of outcomes might be also of interest to marketers.

Customer engagement in the form of likes, comments, and retweets is a key objective to marketers and can easily be tracked at the endorser level. Several studies have already investigated how characteristics of online users are associated with their influence on others (Aral and Walker 2012; Katona et al. 2011; Trusov et al. 2010). However, these studies concentrate on organic word of mouth and voluntary endorsement without monetary incentive (Shi et al. 2014; Toubia and Stephen 2013). Their findings need not generalize to paid endorsement campaigns with monetary incentives. For instance, self-presentation is often a key motive to post online content (Schau and Gilly 2003; Toubia and Stephen 2013), but it is not clear to what extent this holds in paid endorsement and other viral-for-hire campaigns.

This paper aims at filling in this gap in the literature by providing answers to the following questions: (i) how incentive affect endorsers' participation and effectiveness in paid endorsement campaigns, (ii) what types of endorsers are most effective in generating online engagements, and (iii) whether that varies across types of engagements that require different levels of effort from endorsers' followers.

To answer these questions, we collaborated with two vendors on the Chinese retailing site taobao.com, and ran a field experiment on the Chinese microblogging site weibo.com, using one of the largest Chinese paid endorsement platforms, weituitui.com. We exogenously manipulate the pay rate to endorsers and their eligibility to participate. Since the data collected from our experiment are panel count data with sample selection issues, we propose a Poisson lognormal model with sample selection and correlated random effects to analyze what affects endorsers' participation and effectiveness.

Our study produces several intriguing findings. (i) Endorsers are sensitive to losses but not gains, compared to the average reward per task they received in the past. (ii) Observed and unobserved

characteristics of endorsers often have opposite effects on participation and effectiveness. As a result, low potential endorsers may generate high actual engagements due to their high probability to participate, whereas high potential endorsers may generate low actual engagements due to their low probability to participate. (iii) The potential of the same endorser can be different in generating different types of engagement.

This work, as the first attempt to study what affects endorsers' participation and effectiveness in paid endorsement social advertising campaigns, makes the following contributions to the literature. First, it helps marketers understand the role of incentives in such campaigns. Second, it documents a tension between participation and effectiveness, and highlights the difference between potential and actual effectiveness. Third, it suggests that different mechanisms may be driving different types of engagements. Finally, it shows how to deal with sample selection in panel data with repeated observations by combining an exogenous soft eligibility constraint and econometric modeling.

## **2.2 Theoretical Background**

This section discusses motives that may affect endorsers' participation and effectiveness in paid endorsement campaigns, and how financial incentives and three endorsers' characteristics (social media fan base, prior activity level, and community embeddedness) may affect endorsers' participation and effectiveness.

### **2.2.1 Participation**

The literature on survey participation broadly divides the reasons why people participate in surveys or questionnaires into three categories: altruistic reasons (e.g., willingness to help research and civil duty), egoistic reasons (e.g., monetary incentive, opportunity to learn something), and survey-specific reasons (e.g., topical interest, trust in organization) (Singer and Ye 2013). Likewise, in paid endorsement campaigns, the motives of endorsers can be classified into three categories: altruistic (e.g., goodwill to share attractive deals), egoistic (e.g., monetary incentive and self-enhancement), and campaign-specific. In this paper, we only focus on drivers that are relevant to incentive and

endorsers' characteristics. Two such drivers are monetary incentive and self-enhancement (perhaps also goodwill to share attractive deals, which is hard to disentangle with self-enhancement).

*Incentives, reference dependence and loss aversion.* Paid endorsement is predicated on the assumption that financial incentives motivate people to act as endorser. Prospect theory posits that financial rewards motivate people because they are gains or losses compared to some reference point, rather than merely because of the absolute size of the incentive (Kahneman and Tversky 1979; Long and Nasiry). It also posits that people are often more sensitive to losses than gains. The former phenomenon is known as reference dependence and the latter as loss aversion. Both are well-documented in consumer behavior (Greenleaf 1995; Hardie et al. 1993; Kalyanaram and Winer 1995; Lattin and Bucklin 1989). Prior research suggests that the average reward per task that a potential endorser received in the past is a good candidate reference point in paid endorsement campaigns (e.g., Hardie et al. 1993). Both the theory and empirical findings, finally, indicate the presence of decreasing rather than constant returns in how gains and losses affect behavior.

*Self-enhancement.* Theories of self-enhancement suggest that people are motivated to seek positive evaluations from others (Jones 1973). On social media platforms, users' activities are publicly visible to others. This makes social media platforms an ideal place for people to signal their expertise and enhance their social status (Alexandrov et al. 2013; Lovett et al. 2013; Schau and Gilly 2003). In particular, Toubia and Stephen (2013) documented that self-image is the primary motive for most users to contribute content voluntarily to Twitter. Therefore, users with a positive reputation and self-image may be more selective than others in which paid endorsement campaigns to participate. The concern about self-enhancement likely varies with several characteristics of endorsers, as elaborated below.

*Social media fan base* refers to the number of followers that endorsers on social media platforms have. Since the remuneration of endorsers often increases with their number of followers,

endorsers with a larger number of followers might be financially more motivated to participate. However, users with a larger number of followers may derive more self-image related utility (Toubia and Stephen 2013). As a result, they might be more selective about which campaigns to participate in, as broadcasting irrelevant content can hurt their reputation (Barasch and Berger 2014; Bock et al. 2005). Alternatively, it is possible that endorsers with a greater number of followers are more likely to participate regardless of incentive, as they derive more intrinsic and status-related benefits from relaying attractive deals and other interesting content (Toubia and Stephen 2013).

*Prior activity level* refers to the endorsers' past activity intensity on social media and paid endorsement platforms. The more posts a user made on social media, and the more paid endorsement campaigns a user participated in, the less selective the user may be in deciding what to post and what to participate in (Porter and Whitcomb 2003). Therefore, we expect endorsers who posted more and participated more in the past to be more likely to participate in a future campaign.

*Community embeddedness* refers to how long the endorsers have been registered and how many friends they have in the paid endorsement community. Endorsers who are more deeply embedded into the community might be more selective in what campaigns to participate in (Minkler 2012), and more concerned about their status when sharing content in online communities (Schau and Gilly 2003; Toubia and Stephen 2013). Thus, such endorsers may be more selective and less likely to participate in any given endorsement campaign.

### 2.2.2 Effectiveness

The effectiveness of endorsers in generating engagements depends on their level of effort, the trust of their followers in them, the sheer numbers of followers, and the strength of the ties with their followers (Aral and Walker 2014; Chu and Kim 2011; King et al. 2014; Moldoveanu and Baum 2011). We discuss the potential effects of incentive size and the three types of endorser characteristics we study based on how they relate to these four traits.

*Incentive.* In paid endorsement platforms, the remuneration of endorsers is often based solely on their number of followers rather than being contingent on performance. Research on survey response behavior suggest two alternative hypotheses regarding the impact of incentive on response quality when incentives do not depend on performance (Cantor et al. 2008; Singer and Ye 2013). One hypothesis is that, by attracting people who would otherwise not participate, the quality of response declines. The alternative hypothesis is that, by rewarding participants, the quality of responses increases due to feelings of gratitude or obligation. A comprehensive review of studies evaluating the effects of incentive on response quality (e.g., number of questions answered and length of answers) concluded that incentive size almost never had an effect on quality (Singer and Ye 2013). This suggests that in paid endorsement campaigns, the size of incentive need not impact the effectiveness of endorsers. Therefore, we expect little to no effect of incentive on effectiveness.

*Social media fan base.* While the tie strength between users and their contacts decreases with the number of contacts (Burke 2011; Katona et al. 2011; Roberts et al. 2009), a larger fan base implies a larger audience who can potentially engage (Goel et al. 2012). A number of studies have investigated the effect of network size on a user's overall influence, but the results are mixed. Katona et al. (2011) find that the effectiveness of individuals in influencing friends to adopt (register on) a social network site decreases with the total number of their contacts, whereas Yoganarasimhan (2012) finds that a node's overall effectiveness in spreading Youtube videos increases with its network size. One explanation to reconcile these two findings is that the effect of network size depends on the level of effort needed to make a decision. When the required effort is small (e.g., information diffusion), weak ties suffice (Granovetter 1973; Weimann 1983) and the effect of network size is dominated by volume per se, leading to a positive overall effect. On the other hand, when the required effort is large (e.g., product adoption), the need for strong ties (Weenig and Midden 1991; Weimann 1983) make users with larger number of followers connected by weak ties not as persuasive, resulting in a negative overall effect. This implies that the effect of

the number of weak tie followers on comments and retweets might be smaller than that on likes, as comments and retweets require more effort than likes.

*Priority activity level.* Endorsers who posted and participated a lot in the past are less likely to be selective and more likely to be spammers. Numerous posts or endorsements can hurt their reputation, rendering them less trustworthy than those who do not post/endorse as much (Barasch and Berger 2014; Bock et al. 2005). Therefore, endorsers who posted and participated more in the past should be less effective. Note, this implies that endorser characteristics associated with prior activity may have opposite effects on participation and effectiveness.

*Community embeddedness.* Following the argument that endorsers who are embedded into the paid endorsement community tend to be more selective in what to participate, it is likely that their follower will trust their endorsements more. Consequently, endorsers with stronger community embeddedness are expected to be more effective in generating online engagements from their followers. Note, this implies that endorser characteristics associated with community embeddedness may have opposite effects in participation and effectiveness.

## **2.3 Field Experiment**

### **2.3.1 Research Setting**

We conducted a field experiment on weituitui.com, a social advertising platform with more than 40,000 registered endorsers who own accounts on weibo.com. Weituitui.com is a broker website that allows advertisers to recruit endorsers at pre-specified prices for their social media marketing campaigns. An advertiser can initiate a paid endorsement campaign by posting a task describing her needs on weituitui.com. In the task, the advertiser also specifies how much an endorser will be paid, as a linear or step-wise linear function of the endorser's number of followers on weibo.com. To penalize robot followers and inactive followers, weituitui.com uses the number of verified followers to calculate the reward for an endorser. Weituitui.com has an internal algorithm to compute the percentage of verified followers based on how actively an endorser's followers engage

on her past tweets. Similar to other paid endorsement platforms like sandaha.com, weituity.com has several policies in place to make sure that the rewards are not too small to be meaningful and also to encourage endorsers with small numbers of followers to participate. The rewards for endorsers with less than 1000 verified followers are fixed on weituity.com (10-49: 0.1RMB, 50-99: 0.2RMB, 100-499: 0.3RMB, 500-999: 0.5RMB), regardless of the reward structure. Endorsers with less than 10 verified followers are not allowed to participate. The reward for endorsers with more than 1000 verified followers are no less than 0.5RMB.

In a task, the advertiser provides the URL of the target tweet containing the product information. The advertiser can impose some written requirements for the task, such as how long the endorser should keep (i.e., not delete) the retweet on their timeline, and the minimal length of the comment in the retweet. Furthermore, the advertiser can specify who is eligible for the task. Some eligibility restrictions are hard restrictions automated by the platform, such as the allowable day part of participation (e.g., 9am-9pm), while other are soft restrictions attached in the written requirements that need to be manually verified afterwards. If an endorser decides to participate, she needs to retweet the given tweet, fulfill the requirements, and then submit the URL of her retweet. The duration of a task ranges from 3 to 5 days. Once the task ends, the advertiser has 3 days to manually approve or disapprove the submissions, depending on whether the endorser has truly retweeted the given tweet and fulfilled the requirements. All remaining submissions are approved automatically by the platform after the 3-day window. Because of this auto-approval policy, opportunistic endorsers or spammers may submit a random URL even if they haven't retweeted the tweet. For approved tasks, the endorsers are paid and weituity charges a 30% commission fee. That fee is reduced to 15% for an extremely small proportion (0.3%) of VIP endorsers who have spent (rather than earned) more than 1000RMB on weituity.com.

### 2.3.2 Experiment Design

To investigate the effect of incentive on endorsers' willingness to participate and their effectiveness in generating engagements (i.e., likes, comments, and retweets), we exogenously manipulate the

incentive by posting two identical tasks at two different pay rates. We use the linear pricing scheme as it is easier to implement and understand. The two pay rates are 0.0002 RMB (1RMB  $\approx$  0.16USD) and 0.0004 RMB per follower, respectively. The former is the lowest possible and most common rate for linear pricing (i.e., 87% of tasks)<sup>25</sup>, whereas the latter is higher than or equal to 96% of linear rates used on weituitui.com. Figure 2.1 plots the incentive curves for the two pay rates, showing how the number of verified followers maps into the financial rewards at the low and high pay rates. The percentile of endorsers is given on the top of the figure (e.g., 59% endorsers have less than 500 verified followers). Note how rewards at different pay rates differ for only about 23% of endorsers with the most verified followers.

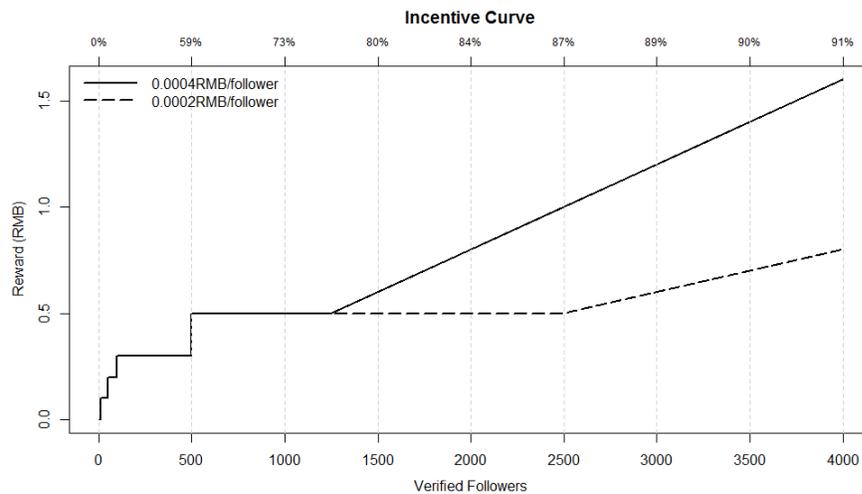


Figure 2.1 The Incentive Curves at Low and High Pay Rates.

To make sure that the two tasks are indeed identical and yet independent with each other, we register a new account on weibo.com and post two identical tweets on the same product at roughly the same time (more precisely, one is posted just seconds ahead of the other). The URLs of the two tweets are then used in the two tasks, respectively. Since the new account has no followers, all the observed engagements on the two tweets come from the paid endorsers and their followers.

<sup>25</sup> Historically, 1% of linearly priced tasks were posted at the rate of 0.00015 RMB per follower, but the minimum rate had been changed to 0.0002 RMB per follower more than two years before we ran the experiment.

To eliminate the potential effects resulting from the order of the two tweets, the pay rates associated with the tasks posted first and second are swapped from time to time.

Since there might be unobserved variables that affect both participation and effectiveness of endorsers, the identification of the effects of endorser characteristics in both stages typically requires an exclusion restriction (Puhani 2000). To that end, we add a soft eligibility restriction in our tasks, such that every endorser on weizuitui.com is only eligible for one of the two identical tasks. “Soft” means that ineligible endorsers can still participate, but will not be paid. This is known as an invitation or encouragement design (Brewer 1976; Duflo and Saez 2003; Powers and Swinton 1984). Specifically, endorsers whose last two digits of their weizuitui IDs (six-digit numbers) are among a certain range are eligible to participate in one task, and those among another range are eligible for the other task. The eligibility restriction is a valid instrument as the last two digits of an endorser’s ID are random and unrelated to her effectiveness.

Table 2.1 Experimental Design

	Pair A	Pair B
Price: 0.0002 RMB/follower	00~24	50~74
Price: 0.0004 RMB/follower	25~49	75~99

Eligible IDs shown in cells were rotated across pay rates and products across weeks.

Our experiment was conducted in 8 different weeks between 2/1/2014 and 4/26/2014. Each week, we posted two groups (pairs) of identical tasks on two products from the same vendor. Accordingly, we divided endorsers into 4 different groups based on their ID (i.e., 00~24, 25~49, 50~74, 75~99), such that any endorser was eligible for only one of the four tasks in that week. The four tasks were posted simultaneously so that they showed up right next to each other. The tasks were rotated over 6 products from 2 vendors on taobao.com. Each task was open for participation for 72 hours. For our experiment, we did not impose any particular effort-related task requirements except for retweeting and liking the tweet. The eligible pay rates for the same endorsers were rotated across weeks. Table 2.1 visualizes the key conditions of our experimental design by showing the four tasks

posted in a given week. Each task pertains to one of two products for which endorsers are promised either a high or low pay rate, and a potential endorser qualifies for only one of the four tasks.

## 2.4 Data

### 2.4.1 Descriptive Statistics

The data is collected and analyzed at the endorser level. We focus on the 8,283 active endorsers who participated in at least one paid endorsement task in the 6 months prior to our experiment. In every task, we record whether each of these active endorsers participates and how many engagements she generates. The number of engagements is collected for each retweeter/endorser using the API provided by weibo.com. Herein, participation means that an endorser has actually retweeted the message. The participation and engagement statistics are summarized in Table 2.2. Excluding one task for which we failed to track the engagement due to a technical issue, the 31 tasks we posted attracted 2,241 participations from 1,016 endorsers.

Table 2.2 Experiment Statistics

Number of weeks	8
Number of products	6
Number of tasks	31
Number of endorsers	8,283
Number of participating endorsers	1,016
Number of participations	2,241
Number of participations from ineligible endorsers	91
Number of <task, endorser> observations	227,608

Some endorsers registered on weizuitui.com in the middle of our experiment

Detailed task-by-task participation and engagement statistics are shown in Table 2.3. There is a clear decline in the number of comments and retweets generated per task over time. The number of likes is also decreasing but not as fast. We stopped the experiment after week 8 due to this saturation effect. To make sure that the tasks did not run out of budget before they were closed for participation, we tried different budgets (100, 200, and 300RMB) in the first four weeks and found that 200RMB was more than enough. Therefore, the budgets for all the tasks in the subsequent

weeks are 200RMB. The distribution of engagement generated by participating endorsers is shown in Table 2.4. Most endorsements do not generate any engagement.

Table 2.3 Participation and Engagement Statistics of Tasks

Week	Pair	Task	Budget	Pay Rate	Product	# Participants	# Followers	Total Engagements		
								likes	comments	retweets
1	1	1	100	4	FT	32	0.6 M	9	13	21
	1	2	100	2	FT	3	1.3 M	8	67	67
	2	3	100	4	HRM	34	0.6 M	17	5	5
	2	4	100	2	HRM	65	0.8 M	11	120	121
2	3	5	200	2	ST	113	1.4 M	8	40	46
	3	6	200	4	ST	109	2.0 M	12	24	26
	4	7	200	2	BL	91	3.6 M	9	14	18
	4	8	200	4	BL	119	2.1 M	15	26	33
3	5	9	300	4	BL	71	2.0 M	8	10	10
	5	10	300	2	BL	84	2.8 M	4	20	34
	6	11	300	4	ST	80	1.1 M	4	49	50
	6	12	300	2	ST	78	2.2 M	6	14	7
4	7	13	300	4	ST	56	0.8 M	3	2	5
	7	14	300	2	ST	66	2.5 M	6	5	3
	8	15	300	4	BL	81	1.4 M	6	25	32
	8	16	300	2	BL	74	1.2 M	4	8	1
5	9	17	200	2	LP	63	1.0 M	13	1	5
	9	18	200	4	LP	77	0.7 M	8	3	0
	10	19	200	2	ER	76	1.8 M	8	1	4
	10	20	200	4	ER	70	1.3 M	0	1	2
6	11	21	200	4	HRM	67	1.3 M	6	4	0
	11	22	200	2	HRM	63	0.5 M	11	5	1
	12	24	200	2	FT	70	1.2 M	2	1	1
7	13	25	200	2	ER	72	1.1 M	8	3	1
	13	26	200	4	ER	70	0.7 M	4	3	1
	14	27	200	2	LP	80	1.5 M	9	8	2
	14	28	200	4	LP	75	1.7 M	2	4	3
8	15	29	200	4	FT	72	1.1 M	6	1	0
	15	30	200	2	FT	79	0.9 M	3	3	2
	16	31	200	4	HRM	81	1.6 M	9	5	6
	16	32	200	2	HRM	70	1.4 M	5	9	0

For pay rate, “2” and “4” represent 0.0002 and 0.0004 RMB/follower, respectively. The number of participants represents the number of endorsers who have retweeted the tweets in the given tasks. The number of followers represents the total follower number of all participated endorsers. The total engagements represent the total number of likes, comments and retweets generated by the participants. The six products used in our experiment are: Heart Rate Meter (HRM), Fitness Tracker (FT), Buddha Statue (ST), Bracelet (BL), Ear Ring (ER), and Lapis Lazuli (LP). HRM and FT were sold by one vendor, and the other four products by another vendor.

Table 2.4 Distribution of Engagements Generated by Individual Endorsers

Type	Distribution									
	0	1	2	3	4	5~10	>10	Mean	SD	Max
likes	2072	145	14	4	4	1	1	0.10	0.50	15
comments	2104	69	21	15	5	16	11	0.22	1.59	34
retweets	2130	50	9	15	7	15	15	0.23	1.71	36

Table 2.5 Description of Independent Variables

Variables	Description
<b>Exclusion Variable</b>	
isEligible	Whether a endorser is eligible for a given task (for selection equation only)
<b>Incentive</b>	
payRate	Pay rate per follower (either 0.0002 or 0.0004 RMB/follower)
actRwd	Actual reward upon approval, net of commission fee
avgRwd	Average reward per task of an endorser in the past
gain	$\text{Max}(0, \text{actRwd} - \text{avgRwd})$
loss	$\text{Max}(0, \text{avgRwd} - \text{actRwd})$
<b>Social Media Fan Base</b>	
followers	Number of followers on weibo.com
verifiedRatio	Percentage of verified followers in all followers
<b>Prior Activity Level</b>	
tweetNum	Number of tweets posted on weibo.com
taskNum	Total number of tasks participated in the past
approvalRate	Percentage of approved tasks in the past
<b>Community Embeddedness</b>	
regDays	Number of days an endorser has registered on weuitui.com (rescaled to [0,1])
friends	Number of friends an endorser has on weuitui.com's internal social network
<b>Other</b>	
group	A dummy accounting for the fixed effects for each of the 16 tasks groups (pairs)
referralRwd	Total reward received through referring others to register on weuitui.com
times	Number of times an endorser has participated in tasks on the same product

We collect data on the characteristics of endorsers by scraping their profiles on weuitui.com, which include their information on both weuitui.com and weibo.com. The information on weuitui.com includes the number of verified followers, the number of tasks participated in, the total amount of reward earned, the total referral income, the number of friends on weuitui's internal social network, and how long ago one registered on weuitui.com. The information on weibo.com includes the number of followers and the number of tweets (including retweets). In Table 2.5, we summarize the independent variables used for our analysis in six different categories. We focus on those variables that advertisers can set or observe and hence use for targeting. These variables fall into four categories: incentive, social media fan base, prior activity level, and community embeddedness. For incentive, in addition to pay rate, we also compute the actual reward an endorser will receive upon approval, which allows us to better account for the special pricing scheme showing in Figure 1. Note that the actual reward for ineligible participants will be zero regardless of their number of

verified followers.<sup>26</sup> As we discussed earlier, the effect of incentive may depend on some reference level. We choose the average reward per task in the past as the reference point and then derive the gain and loss for each endorser. The variables in the “Other” category, such as “referralRwd” and “times”, are specific to the platform and our experimental design, and are used merely as controls. They are not of substantive interest.

Table 2.6 Key Statistics on Independent Variables

Variables	Entire Dataset					Subset of Participants				
	Mean	Median	Min	Max	SD	Mean	Median	Min	Max	SD
isEligible	0.25	0.00	0.00	1.00	0.43	0.96	1.00	0.00	1.00	0.20
payRate	2.97	2.00	2.00	4.00	1.00	2.98	2.00	2.00	4.00	1.00
log(actRwd)	-5.53	-6.91	-6.91	5.37	2.44	-1.14	-1.05	-6.91	4.29	1.54
log(gain)	-6.51	-6.91	-6.91	5.27	1.30	-4.68	-6.91	-6.91	3.69	2.71
log(loss)	-6.18	-6.91	-6.91	3.27	1.94	-4.44	-6.91	-6.91	2.74	2.71
log(avgRwd)	-0.92	-0.92	-6.91	3.48	0.97	-0.72	-0.84	-2.80	3.48	0.72
log(followers)	6.93	6.84	2.48	15.42	2.02	7.88	7.76	2.83	14.51	1.97
verifiedRatio	0.44	0.46	0.00	1.00	0.24	0.45	0.46	0.00	1.00	0.25
log(tweetNum)	5.90	5.99	0.00	11.26	1.75	6.57	6.66	0.00	11.26	1.61
log(taskNum)	2.48	2.30	0.00	8.70	1.87	4.54	4.76	0.00	8.70	1.81
approvalRate	0.74	0.82	0.00	1.00	0.28	0.83	0.87	0.00	1.00	0.16
regDays	0.21	0.18	0.00	1.00	0.15	0.21	0.17	0.00	1.00	0.19
log(friends)	0.38	0.00	0.00	2.77	0.71	0.60	0.00	0.00	2.77	0.90
log(referralRwd)	-5.54	-6.91	-6.91	7.17	2.94	-4.69	-6.91	-6.91	5.45	3.55
times	0.02	0.00	0.00	4.00	0.15	0.09	0.00	0.00	4.00	0.32

Table 2.7 Correlation between Independent Variables

isEligible	1.00																
payRate	0.00	1.00															
log(actRwd)	0.98	0.01	1.00														
log(gain)	0.53	0.05	0.60	1.00													
log(loss)	0.65	-0.02	0.60	-0.11	1.00												
log(avgRwd)	0.00	0.00	0.04	-0.10	0.21	1.00											
log(followers)	0.00	0.00	0.08	0.16	-0.07	0.36	1.00										
verifiedRatio	0.00	0.00	0.01	0.04	-0.04	-0.08	-0.34	1.00									
log(tweetNum)	0.00	0.00	0.04	0.07	-0.03	0.17	0.49	-0.28	1.00								
log(taskNum)	0.00	0.00	0.02	-0.04	0.13	0.29	0.21	-0.04	0.22	1.00							
approvalRate	0.00	0.00	0.01	0.07	0.06	0.00	0.11	-0.02	0.10	0.21	1.00						
regDays	0.00	0.00	0.02	0.01	0.01	0.07	0.23	-0.10	0.27	0.44	0.09	1.00					
log(friends)	0.00	0.00	0.00	-0.04	0.07	0.10	0.00	-0.02	0.07	0.40	0.02	0.22	1.00				
log(referralRwd)	0.00	0.00	0.00	-0.03	0.07	0.14	0.07	-0.04	0.10	0.42	0.09	0.29	0.36	1.00			
times	-0.04	0.00	-0.04	-0.02	-0.03	0.03	0.07	0.00	0.06	0.15	0.05	0.00	0.04	0.04	1.00		

<sup>26</sup> We have also tried an alternative version of reward which does not distinguish between eligible and ineligible endorsers in computing expected rewards. That is, even ineligible endorsers can have non-zero rewards. In our later analyses, we find that this alternative coding produces very similar findings but worse model fit.

The summary statistics of the independent variables are shown in Table 2.6. The characteristics of participating endorsers are clearly different from those of the whole population, which is evidence of self-selection. The correlations among the independent variables are shown in Table 2.7. Except for the expected correlations with variables representing incentive, the two manipulated variables “isEligible” and “payRate” have zero correlation with other variables, indicating effective randomization. isEligible is correlated with incentive because only eligible endorsers can have positive actual reward.

#### 2.4.2 Model-Free Analysis of Manipulation Effects

To provide some intuition regarding how the manipulations affect the participation and effectiveness of endorsers, we compare the participation rates and generated engagements between eligible vs. ineligible and between high-pay rate vs. low-pay rate endorsers. In addition, given that eligibility and pay rate may affect the effort level of endorsers, we also compare the effort levels of endorsers in different treatment groups. In paid retweeting campaigns similar to ours, the only place where endorsers can show differentiated efforts lies in the composition of the comment included in the retweet, if any. Two metrics that reflect the effort level of an endorser in composing a comment is the length of comment (namely the number of words<sup>27</sup>) and the use of emoji (yes or no). The former metric is commonly used to measure the effort level of respondents (Singer and Ye 2013).

Table 2.8 contrasts the average participation rates, effort levels, and engagements of eligible vs. ineligible (high-pay vs. low-pay) endorsers. The *p*-values from ANOVA test are provided. As expected, eligibility has a strong effect on participation but no effect on effort or engagement level. Therefore, eligibility is indeed a valid instrument, as expected. However, the effects of pay rate may be a bit surprising, as there are no significant differences between high and low pay rate except for the usage of emoji. The most likely explanation is that the payment for 77% of endorsers is not affected by high vs. low pay rates, due to weizitui’s constraint on the reward structure, and that the

---

<sup>27</sup> In Chinese, one word is represented by one character, so the number of characters is the same as the number of words.

effect of high vs. low pay rate on incentives received varies as a function of the number of verified followers for the remaining 23% (Figure 2.1).

Table 2.8 Effects of Manipulated Variables

Manipulated Variables		Participation rate	Effort Levels in Retweets		Engagements		
			# words attached	emoji	Likes	Comments	Retweets
<b>Eligibility</b>	Eligible endorsers	3.78%	16.54 (11.28)	9.77%	0.10 (0.51)	0.22 (1.57)	0.22 (1.70)
	Ineligible endorsers	0.05%	15.34 (12.09)	5.49%	0.09 (0.28)	0.33 (1.89)	0.42 (1.96)
	ANOVA (p-value)	<0.001	0.32	0.18	0.82	0.50	0.28
<b>Pay Rate</b>	Low-paid endorsers	0.97%	16.59 (11.01)	8.11%	0.10 (0.44)	0.28 (2.02)	0.27 (2.12)
	High-paid endorsers	1.00%	16.39 (11.63)	11.15%	0.10 (0.56)	0.16 (0.93)	0.18 (1.12)
	ANOVA (p-value)	0.61	0.68	0.01	0.98	0.08	0.19

Reported values are means, with standard deviation in brackets. Both effort levels and engagements are conditional on participation.

## 2.5 Model

The data analysis presents two challenges. First, engagement is observed only for those endorsers who participate in the task, and the effectiveness of participants may not be representative of the whole population. This is commonly known as the sample selection problem (Heckman 1979). Second, an endorser can participate in more than one task and the resulting observations on the same endorser may not be independent. While both the sample selection and repeated observation problems are common in the literature and can be addressed effectively when they appear separately, little has been done to address both problems jointly, especially when the dependent variable is counts. We propose a model to deal with both problems. We first present our approach to model participation and potential effectiveness of endorsers jointly in Section 2.5.1, and then discuss its connection to existing models in Section 2.5.2. We elaborate on how to compute some effects of substantive interest in Section 2.5.3.

### 2.5.1 Sample Selection Model with Correlated Random Effects

We model likes, comments, and retweets separately. For each of these outcomes, there are two equations in our model: the selection or participation equation captures what affects an endorser's

participation, and the outcome equation captures what affects an endorser's potential effectiveness in generating engagement. The potential effectiveness is not conditional on participation, which allows us to gain insights on the entire population of endorsers, not only on those who participated. We use boldface letters to represent vectors and matrices. For notational compactness, we use row vectors throughout this chapter.

Following the standard sample selection model (Greene 2009; Heckman 1979), we use a probit model for the selection equation. Letting the variable  $z_{it}$  indicate whether endorser  $i$  participates in task  $t$ , the participation decision is given by

$$z_{it} = \mathbf{1}(\boldsymbol{\alpha}\mathbf{w}'_{it} + \delta u_i + \xi_{it} > 0) \quad (2.1)$$

where  $\mathbf{w}_{it}$  includes an intercept and the sets of variables that affects the participation decision of endorser  $i$  in task  $t$ . The variables in  $\mathbf{w}_{it}$  include characteristics of endorser  $i$ , characteristics of task  $t$ , the characteristics specific to the endorser-task dyad, and the exclusion variable (see Table 2.5). They also include 15 dummy variables for each pair of identical tasks posted that vary only on pay rate (the intercept captures the sixteenth pair). These dummies absorb any task-specific effect apart from pay rate, including characteristics of the product featured, characteristics of our post on weizuitui, and temporal shocks. Note that the two paired tasks share the same fixed effect as they are identical except for pay rate which is controlled separately. The random terms  $u_i \sim N(0,1)$  and  $\xi_{it} \sim N(0,1)$  capture endorser and endorser-task level unobserved characteristics that affect the participation decision, respectively. The selection equation given above is a probit model with random effects (Butler and Moffitt 1982).

Since the engagements (including likes, comments, and retweets) are all counts, and since the data feature both overdispersion (see Table 2.4) and repeated observations, we use a Poisson lognormal model with random effects for the outcome equation. Let  $y_{it}^*$  be the potential outcome (i.e., the number of likes, comments, or retweets) of endorser  $i$  on task  $t$ . The outcome equation is given by

$$E[y_{it}^* | \mathbf{x}_{it}, \varepsilon_i, \varepsilon_{it}] = \lambda_{it} = \exp(\boldsymbol{\beta} \mathbf{x}_{it}' + \sigma \varepsilon_i + \gamma \varepsilon_{it}) \quad (2.2)$$

where  $\mathbf{x}_{it}$  includes an intercept and the set of variables that affects the potential engagement generated by endorser  $i$  for task  $t$ . The only difference between  $\mathbf{x}_{it}$  and  $\mathbf{w}_{it}$  is that only  $\mathbf{w}_{it}$  includes the exclusion variable. Our outcome equation accounts for two levels of heterogeneity:  $\varepsilon_i \sim N(0,1)$  and  $\varepsilon_{it} \sim N(0,1)$  capture the effect of endorser and endorser-task level unobserved characteristics, respectively. When  $\sigma = 0$ , the above model simplifies to the Poisson lognormal model (Greene 2009), which often yields similar estimates to the Negative Binomial model. In our data, we find that the above model with two levels of heterogeneity fit the data substantially better than the zero-inflated Poisson and the Negative Binomial models.

The error terms in the selection and outcome equations need not be independent. Specifically, the endorser-level unobserved characteristics that affect selection or participation may also affect outcomes, and so may endorser-task level unobserved characteristics. As a result, we further assume that the endorser and endorser-task level error terms are bivariate normally distributed, with a correlation of  $\rho$  and  $\tau$ , respectively.

$$\begin{pmatrix} u_i \\ \varepsilon_i \end{pmatrix} \sim N \left( \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right), \begin{pmatrix} \xi_{it} \\ \varepsilon_{it} \end{pmatrix} \sim N \left( \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \tau \\ \tau & 1 \end{pmatrix} \right).$$

Compared to the existing sample selection models (Greene 2009; Heckman 1979; Winkelmann 1998), our model not only takes into account random effects, but also allows the random effects to be correlated. Letting  $T_i$  be the number of tasks endorser  $i$  can potentially participate, the likelihood of all observations on endorser  $i$  can be written as

$$\begin{aligned} L_i &= P(y_{i1}^*, \dots, y_{iT_i}^*; z_{i1}, \dots, z_{iT_i} | \mathbf{x}_{i1}, \dots, \mathbf{x}_{iT_i}, \mathbf{w}_{i1}, \dots, \mathbf{w}_{iT_i}) \\ &= \int_{-\infty}^{\infty} \phi(\varepsilon_i) d\varepsilon_i \int_{-\infty}^{\infty} f(u_i | \varepsilon_i) du_i \prod_{t=1}^{T_i} \int_{-\infty}^{\infty} P(y_{it}^* | \mathbf{x}_{it}, \varepsilon_i, \varepsilon_{it})^{z_{it}} P(z_{it} | \mathbf{w}_{it}, u_i, \varepsilon_{it}) \phi(\varepsilon_{it}) d\varepsilon_{it} \quad (2.3) \end{aligned}$$

where  $P(y_{it}^* | \mathbf{x}_{it}, \varepsilon_i, \varepsilon_{it}) = \frac{\lambda_{it}^{y_{it}^*} e^{-\lambda_{it}}}{y_{it}!}$ , as given by the conditional Poisson distribution. In the likelihood function,  $P(y_{it}^* | \mathbf{x}_{it}, \varepsilon_i, \varepsilon_{it})$  only factors in when  $z_{it} = 1$ , as  $y_{it}^*$  is only observed for participating endorsers. The conditional density  $f(u_i | \varepsilon_i)$  is derived from the bivariate normal distribution. The likelihood of our model does not have a closed-form representation. However, it can be numerically approximated by the Gauss-Hermite quadrature method (Abramowitz and Stegun 1972; Greene 2009). Appendix 2.1 provides more details on the likelihood derivation and parameter estimation. Our model can deal with outcomes following other distributions by changing the distributional assumption on  $P(y_{it}^* | \mathbf{x}_{it}, \varepsilon_i, \varepsilon_{it})$ .

The random terms for different types of engagement may be correlated. However, since the three types of engagements have exactly the same set of regressors in our analysis, estimating the equations for different types of engagements independently, as if there are no correlation across engagements, will give identical estimates (Kruskal 1968).

## 2.5.2 Connections with Existing Models

To convey the connection of the proposed model with existing models, Table 2.9 summarizes potential nested models based on the specification of endorser and endorser-task level random terms. For both endorser and endorser-task random terms, Table 2.9 considers three possibilities: the random term is not specified in either the selection or outcome equation, it is specified in both equations but not correlated, and it is specified in both equations and correlated. For simplicity, Table 2.9 ignores cases in which the random term is specified in one equation but not the other. To allow for unrestricted correlation between the random terms across equations, all the terms are assumed to be normally distributed. In count models, normal error (as in the Poisson Lognormal model) often yields very similar results with Exponential-Gamma error (as in the Negative Binomial model). The models in Table 2.9 are named based on the outcome equation and the correlation(s). The selection equation is probit or probit with random effects by default. Existing models for panel count data with sample selection are rather restricted. The proposed model nests them.

Table 2.9 Models for Panel Count Data with Sample Selection

Endorser level ( $\epsilon_i$ ) \ Endorder-task level ( $\epsilon_{it}$ )	No	Uncorrelated	Correlated
No	Poisson	PLN (Greene 2009)	PLN with SS (Greene 2009)
Uncorrelated	Poisson RE (Hall 2000)		
Correlated			PLN with SS and CRE (Proposed)

PLN: Poisson Lognormal model; RE: random effects; SS: sample selection, namely with correlation on endorser-task level random terms; CRE: correlated random effects, namely with correlation on endorser level random terms. Empty cells represent models that are missing in the literature.

### 2.5.3 Relative Partial Effects on Potential and Actual Outcome

One question of interest to marketers is how much the mean potential outcome  $E[y_{it}^* | x_{it}]$  varies with respect to the changes in endorsers' characteristics  $x_{it}$ . Integrating out  $\epsilon_i$  and  $\epsilon_{it}$  in Equation (2.2) yields

$$E[y_{it}^* | x_{it}] = E_{\epsilon_i} E_{\epsilon_{it}} [E[y_{it}^* | x_{it}, \epsilon_i, \epsilon_{it}]] = \exp\left(\beta x_{it}' + \frac{\sigma^2 + \gamma^2}{2}\right) \quad (2.4)$$

Therefore, the relative partial effects of  $x_{it}$  on the mean potential outcome is simply  $\frac{\partial \log E[y_{it}^* | x_{it}]}{\partial x_{it}} =$

$\beta$ . The absolute partial effects of  $x_{it}$  on  $E[y_{it}^* | x_{it}]$  is  $\exp\left(\beta x_{it}' + \frac{\sigma^2 + \gamma^2}{2}\right) \beta$ . We focus on the relative partial effects as the absolute partial effects are quite sensitive to outliers in the data due to the exponential term.

Advertisers are interested in how the characteristics of endorsers impact not only the potential outcome but also the actual or observable outcome. If an endorser chooses not to participate, the engagements generated would be zero. Therefore, the relationship between the actual outcome  $y_{it}$  and the potential outcome  $y_{it}^*$  can be written as

$$y_{it} = z_{it} y_{it}^* \quad (2.5)$$

Given that  $z_{it}$  and  $y_{it}^*$  are independent conditional on  $x_{it}, w_{it}, u_i, \epsilon_i, \epsilon_{it}$ , the conditional mean outcome can be written as

$$E[y_{it} | \mathbf{x}_{it}, \mathbf{w}_{it}, u_i, \varepsilon_i, \epsilon_{it}] = P(z_{it} = 1 | \mathbf{w}_{it}, u_i, \epsilon_{it}) E[y_{it}^* | \mathbf{x}_{it}, \varepsilon_i, \epsilon_{it}] \quad (2.6)$$

The unconditional (i.e., not conditional on any unobserved variable) mean outcome can be obtained by integrating out  $u_i$ ,  $\varepsilon_i$  and  $\epsilon_{it}$  in Equation (2.6) (see Appendix 2.2 for details).

$$E[y_{it} | \mathbf{x}_{it}, \mathbf{w}_{it}] = E_{u_i} E_{\varepsilon_i | u_i} E_{\epsilon_{it}} [E[y_{it} | \mathbf{x}_{it}, \mathbf{w}_{it}, u_i, \varepsilon_i, \epsilon_{it}]] \quad (2.7)$$

The relative partial effect of a variable  $s_{it}$  on the mean outcome can be written as

$$g_{s_{it}} = \frac{\partial \log E[y_{it} | \mathbf{x}_{it}, \mathbf{w}_{it}]}{\partial s_{it}} = c_{it} \alpha_s + \beta_s \quad (2.8)$$

where  $\alpha_s$  represents the corresponding coefficient in  $\boldsymbol{\alpha}$  if  $s_{it}$  belongs to  $\mathbf{w}_{it}$ , otherwise 0. Similarly,  $\beta_s$  represents the corresponding coefficient in  $\boldsymbol{\beta}$  if  $s_{it}$  belongs to  $\mathbf{x}_{it}$ , otherwise 0. The functional form of coefficient  $c_{it}$ , which is always positive, is given in Appendix 2.3. The standard errors of the relative partial effects can be estimated using the delta method (see Appendix 2.3). It can be seen from Equation (2.8) that, if a variable only affects the participation equation, then its directional effect on actual outcome is consistent with its directional effect on participation. If a variable only affects the outcome equation, then its impact on potential and actual outcomes are in the same direction. However, for variables appearing in both participation and outcome equations, their effects on potential and actual outcomes might have opposite signs.

## 2.6 Results

### 2.6.1 Selection of Incentive Variables

The incentive of endorsers can be represented in three different ways. The first is to use the pay rate which are exogenously manipulated. However, due to the policies of weuitui.com, 77% of endorsers are indifferent between the two pay rates. The second way is to use the actual reward upon approval, which can account for institutional details. The third is to separate rewards into gains and losses, as posited by prospect theory. An additional consideration is that rewards, gains

and losses may have constant or decreasing returns to scale, the latter of which can be accounted for by a log transformation (Hardie et al 1993).

Rather than assuming a priori one of these representations, we choose the one that best fits the data. Since the incentive of endorsers is not contingent on performance, we expect little impact of incentive on outcome. Therefore, we base our choice of representation on how well each fits the participation data. Table 2.10 reports the model fit of different incentive representations. The logged gain-loss representation fits the data best on all three model fitness metrics, as expected from prospect theory. Therefore, in our following analysis, we represent incentive in a gain-loss framework with log transformation. Note, the logged gain-loss representation fits the data best also if we take outcomes into consideration.

Table 2.10 Selection of Incentive Variables

		Linear		Log	
	PayRate	Reward	GainLoss	Reward	GainLoss
-2*LL	13315.9	13313.0	13289.0	13269.3	13239.4
AIC	13371.9	13369.0	13349.0	13325.3	13299.4
BIC	13661.3	13658.4	13659.0	13614.7	13609.5

## 2.6.2 Main Results

Table 2.11 reports the parameter estimates using our main model in Section 2.5. To ease comparison, the estimates for the participation equations for each of the three outcomes are presented side by side, followed by the estimates for the three outcome equations. All the heterogeneity and correlation parameters, as well as model fitness metrics, are shown in the outcome column. We discuss our findings on the independent variables category by category.

*Incentive.* The effect of losses is greater than that of gains, which is consistent with previous evidence of loss aversion (Hardie et al. 1993; Kalyanaram and Winer 1995). Gains have no significant effect on participation, which is consistent with previous evidence that additional incentives do not impact performance when workers already feel adequately remunerated (Cohn

et al. 2015). Our findings suggest that improving participation rates through increasing incentive can be inefficient. To improve participation rates of paid endorsement campaigns, it may be better to provide an incentive just comparable to endorsers' past rewards and then focus on providing non-monetary incentives. The finding that neither gains nor losses have an effect on outcome is consistent with previous findings that incentives need not affect performance unless they are contingent on performance (Singer and Ye 2013).

Table 2.11 Parameter Estimates for Different Types of Engagements (Log GainLoss)

	Selection			Outcome		
	likes	comments	retweets	likes	comments	retweets
<b>Exclusion Variable</b>						
isEligible	2.724***	2.723***	2.720***			
<b>Incentive</b>						
log(gain)	0.007	0.008	0.009	-0.039	-0.011	-0.061
log(loss)	-0.073***	-0.074***	-0.073***	0.022	0.146	0.095
log(avgRwd)	0.034	0.036	0.037	-0.092	0.174	0.333
<b>Social Media Fanbase</b>						
log(followers)	0.083***	0.077***	0.079***	0.221.	-0.039	-0.195
verifiedRatio	0.256*	0.297**	0.283**	0.652	1.105	0.989
<b>Prior Activity Level</b>						
log(tweetNum)	0.035.	0.045*	0.039.	0.001	-0.478***	-0.252*
log(taskNum)	0.555***	0.564***	0.561***	-0.379***	-0.483***	-0.570***
approvalRate	0.063	0.051	0.063	0.704	1.708.	-0.300
<b>Community Embeddedness</b>						
regDays	-4.375***	-4.410***	-4.358***	0.102	2.826*	3.258**
log(friends)	-0.088*	-0.084*	-0.086*	0.420**	0.131	0.368
<b>Others</b>						
log(referralRwd)	-0.015.	-0.017.	-0.017.	0.023	0.003	-0.095
times	-0.142**	-0.142**	-0.147**	-0.285	-0.102	0.017
<b>Heterogeneity</b>						
$\delta$ (selection)				1.310***	1.315***	1.314***
$\sigma$ (outcome)				1.609***	2.190***	2.693***
$\gamma$ (outcome)				0.097	1.131***	1.371**
<b>Correlation</b>						
$\rho$ (endorser)				-0.216***	-0.253***	-0.247***
$\tau$ (endorser-task)				0.008	0.199	0.312.
<b>Fitness</b>						
Log Likelihood				-7229.7	-7304.5	-7216.5
AIC				14583.4	14733.0	14557.0
BIC				15224.2	15373.8	15197.8

Significance codes: "." for  $p < 10\%$ , "\*" for  $p < 0.05$ , "\*\*\*" for  $p < 0.01$ , and "\*\*\*\*" for  $p < 0.001$ . For compactness, the intercept and the coefficients on the dummy variable "taskDummy" are omitted. The level of efforts required for an engagement: like < comment < retweet.

We also estimate a set of models using the exogenously manipulated pay rate to represent incentive (Table A2.1 in Appendix 2.5). We find that pay rate has no effect on participation, even though the higher pay rate used in our experiment exceeds or equals to 96% of linear prices ever

used on *weitungui*. This is consistent with the model-free evidence reported in Section 2.4.2.<sup>28</sup> Using pay rate instead of gains and losses does not affect the coefficients of other variables much. Using other types of incentive representations listed in Table 2.10 yields similar findings (Tables A2.1-2.4 in Appendix 2.5). In the rest of our discussion, we focus on the results reported in Table 2.11 which give the best model fits.

*Social media fan base.* Endorsers with more followers and a higher verified ratio are more likely to participate. One very plausible explanation is that endorsers with a greater number of verified followers derive greater status enhancement from relaying attractive deals than endorsers with fewer verified followers (Toubia and Stephen 2013). An alternative explanation is that, since followers and verified ratio determine the number of verified followers and hence affect the reward of endorsers, this finding indicates that those who are paid more are more likely to participate. This alternative explanation, however, is at odds with the finding that gains have no effect on participation.

Turning our attention to the outcome equation, we see that the number of followers has a marginally significant effect on likes, but not on comments or retweets. This finding, though a bit weak (but rather robust in our analyses), is consistent with our conjecture that the effect of followers may be smaller for forms of engagement that are more effortful. The reason is that higher levels of engagement are facilitated by strong ties, whereas the tie strength between endorsers and their fans decreases with the number of followers (Burke 2011; Katona et al. 2011; Roberts et al. 2009). This finding is consistent with previous findings that network size has a positive effect on overall influence in low-effort behaviors (e.g., (Yoganarasimhan 2012)).

*Prior activity level.* Endorsers who tweeted more on microblogs and who participated in more tasks in the past are more likely to participate in our tasks, yet are less effective in generating

---

<sup>28</sup> Additional analysis shows that pay rate has no effect on participation even for endorsers whose incentives are sensitive to pay rate. We conducted this analysis by interacting the pay rate with a dummy variable indicating whether an endorser is sensitive to pay rate, i.e., whether the endorser has more than 1250 verified followers. Neither the pay rate nor the interaction term have a significant effect.

engagements. This finding confirms our earlier conjecture that endorsers who are less selective tend to be less effective. Specifically, posting irrelevant or unsound content (e.g., ads) too often can hurt users' reputation in online communities (Barasch and Berger 2014; Bock et al. 2005), and hence their effectiveness.

The variable "approvalRate", defined as the percentage of approved tasks in the past, is a metric of endorsers' diligence. To our surprise, we find no significant effect of approval rate in the outcome equation, except for a marginally significant effect on comments. This suggests that approval rate might not be an ideal indicator of quality. Endorsers with higher approval rate might just be more skillful in fulfilling the requirements of advertisers. Given that approval rate has been widely used as a metric to judge the quality of workers in crowdsourcing services such as Amazon Mechanical Turk (Ipeirotis et al. 2010; Paolacci et al. 2010), this finding suggests that the construct validity of approval rate as a metric of quality may need more thorough investigation, at least in the context of paid endorsement campaigns.

*Community embeddedness.* Endorsers who have registered for a longer time and who have more friends on weizhizui's internal social network are less likely to participate in a task, but more likely to generate certain types of engagements (comments and retweets for the number of days since registration, and likes for the number of friends). The opposite effects of these two variables on participation vs. effectiveness suggests that endorsers who are more embedded and respected within the community tend to be more selective, rendering them more effective in generating engagements.

*Unobserved endorser traits.* The negative correlation  $\rho$  is of particular note. It indicates that the opposite effects on participation vs. effectiveness extend to unobserved endorser characteristics.

*Engagement types.* Among the three types of engagements, "likes" require the least effort as they do not involve any typing, and "retweet" require the most effort as they involve both commenting and sharing. The same is true for the three types of engagements on Facebook: likes, comments

and shares (similar to retweets). Facebook assigns the largest weight to shares and the least weight to likes in their EdgeRank algorithm.<sup>29</sup>

The effects of variables in the outcome equations are often different for relatively high-effort engagement through comments and retweets versus low-effort engagement through likes. For example, it's harder for endorsers who tweet and endorse a lot to generate high-effort engagement (i.e., comments and retweets) than low-effort engagement (i.e., likes). The reverse is true for having been registered for a long time. These findings suggest that having been selective in the past and being a long-time endorser are associated with being more effective, not just in general but especially so for high-quality engagement.

### 2.6.3 Robustness

*Identification.* The validity of the exclusion restriction is critical to our analysis. One conceivable concern might be that, though the eligibility constraint was assigned randomly and independently of any endorser trait, the imposed (in)eligibility changed the endorsement behavior (e.g., effort level) of the endorser and hence affected her effectiveness indirectly. For example, ineligible endorsers may exert stronger effort than eligible endorsers in order to be approved, or exert lower effort given that they have lower faith in actually getting paid. If that is true, eligibility might not be truly exclusive. However, simple ANOVA analyses reported in Table 2.8 shows that eligibility has no effect on effort. More sophisticated multivariate analyses show that this conclusion is robust to controlling for the independent variables entering the outcome equation (Table A2.5 in the Appendix 2.6). The concern that eligibility might have affected the effort level and hence the effectiveness is not supported by the data.

Another concern about the validity of our analysis and findings is that the tasks may have interfered with each other, even though we used a unique tweet for each task. A first cause for concern is that there might be overlap between two endorsers' set of followers. If a follower is exposed to more

---

<sup>29</sup> <http://www.socialbakers.com/blog/1304-understanding-increasing-facebook-edgerank>

than one endorser's retweets, there might be an attribution problem on the engagement by that follower. This type of interference is unlikely in our data since, unlike Twitter, Weibo allows users to engage (like, comment, retweet) on different retweets of the same tweet separately. If needed, the follower can like/comment/retweet all the retweets from different endorsers. Moreover, even if the follower decided to interact with only one of the retweets on the same product, the retweet the follower actually engages on is likely to be from the endorser who has the primary effect on the follower's decision.

A second cause of concern about interference is that, as we posted multiple tasks on the same product over time, there may have been a saturation effect if a follower saw the same product endorsed multiple times. While such interference may indeed have depressed the average effectiveness, there is no compelling reason to believe it would bias our estimates in the selection and outcome equations in opposite directions in a systematic manner. In other words, our main findings that participation and effectiveness are often at odds is not likely to be driven by the saturation effect. Moreover, the task level dummy variable and the control variable "times" already accounted for any main effect of task-level and endorser-level saturation on effectiveness.

A final concern is that, among all participating endorsers, 16 participated in two tasks in a same group (i.e., two identical tasks on the same product at different pay rates), leading to a potential attribution problem between the two retweets on the same product retweeted by the same endorser. However, in the corresponding 16 (endorsers) \* 2 (tasks) \* 3 (types of engagements) observations, only 6 have non-zero engagements (max is 3 and median is 1). The attribution issues on such a small number of observations with such low engagements are unlikely to bias our estimates substantially.

*Robustness to model complexity.* In Table 2.11,  $\gamma$  is insignificant for likes and  $\tau$  is insignificant for likes and comments (the significance for retweets is also only marginal), which might be a signal for over-specification. To examine whether our findings are an artifact of over-specification, we force  $\tau$  (and  $\gamma$ ) to zero and re-estimate the parameters. The results in these simplified models are

very similar to the full model (see Table A2.6 in Appendix 2.6). Since we are ignorant a priori about which of the parameters are significant, for the sake of full information disclosure, we report the results from the full model as main results. We further tried some additional simplified models, including removing random effects, removing dyadic heterogeneity, and removing all correlations (i.e., estimating selection and outcome equation models separately), and find that the findings of substantive interest are highly robust.

*Robustness to outliers.* Among the 31 tasks, task 2 seems to be an outlier, with only 3 valid participations. This is because an endorser with half a million verified followers participated in this task soon after the task was posted and exhausted the budget of the task, which prevented other endorsers from participating. In fact, the other three tasks in the first week may also suffer somehow from this problem, due to the relatively lower budget. However, for tasks after the first week, this should not be an issue, as the budget was either double or triple that in week 1. To assess whether our findings are driven by the potential outlier task(s) in week 1, we repeated our analysis first after removing task 2 and then after removing tasks 1-4 (all tasks in week 1). The estimates and findings of substantive interest are robust, except that the confidence bounds widen somewhat due to the smaller number of observations (Table A2.7 in Appendix 2.6).

## **2.7 Implications for Program Design**

### **2.7.1 Influencing Endorsers by Redesigning Tasks**

The results in Table 2.11 show that endorsers who are responsive to the campaigns are often less effective, whereas effective endorsers are often less responsive. To better understand the tension between responsiveness and effectiveness, we grouped endorsers into four cells in Table 2.12 based on their predicted responsiveness and effectiveness (Appendix 2.4 presents expressions for predicted values). Table 2.12 labels an endorser effective (responsive) if her predicted potential to generate engagements (predicted probability to participate) is above the mean of the data set. The percentages are first computed for individual tasks and then averaged over all tasks. While making

predictions, we assume that every endorser is eligible for every task and the incentive for an endorser is her average reward per task in the past. This rules out the effects of the manipulated eligibility and incentive, allowing us to focus on the effects of endorsers' characteristics. The results are very similar if we assume that every endorser is paid at either the lower or higher rate.

Table 2.12 Distribution of Endorsers

	Likes		Comments		Retweets	
	Effective	Ineffective	Effective	Ineffective	Effective	Ineffective
Responsive	3.5%	26.7%	1.1%	29.0%	0.5%	29.7%
Unresponsive	32.6%	37.2%	25.6%	44.3%	30.2%	39.6%

Only a very small percentage of endorsers are both effective and responsive. This is especially so for higher-effort engagements, i.e., comments and retweets. To improve the effectiveness of paid endorsement campaigns, advertisers may want to find ways to attract endorsers who are effective but unresponsive (e.g., endorsers who have registered for a long time and have many friends on weuitui.com). For example, to attract selective and effective endorsers, advertisers may want to experiment with designing ads that are less likely to hurt an endorser's reputation (e.g., native ads that look like organic tweets). They may also want to experiment with lowering the task requirements and offering tasks exclusively to endorsers who have registered for a long time and who have many friends on weuitui. This can be implemented by the written eligibility restrictions in the tasks.

In addition, advertisers may want to seek ways to improve the impact of responsive but ineffective endorsers (e.g., endorsers who tweet and endorse a lot). For example, if the expected participants are mostly those who are responsive but not effective, advertisers can increase the effort-related requirements in the tasks such as the minimal number of words and emojis in retweets, and the minimal number of people to be mentioned while retweeting.

In practice, it's possible to "attract" effective but unresponsive endorsers and "enforce" responsive but ineffective endorsers at the same time by offering them different versions of tasks exclusively.

For instance, if the objective of advertisers is to generate comments or retweets, advertisers can offer an “attract” task to endorsers who post few tweets and participate in few tasks, and an “enforce” task to endorsers who post a lot and participate a lot.

### 2.7.2 Boosting Potential vs. Actual Engagements by Targeting

The tension between responsiveness and effectiveness further invites analysis of the relative partial effects of the independent variables on the actual vs. potential engagements. For simplicity, we call effects on the actual outcomes the “total effects” through both participation and potential outcomes. To generate a large number of engagements, an endorser needs to not only have high a potential to generate engagements, but also to actually participate in the campaign. The total effect of a variable on actual engagement can be computed using Equation (2.8), which represents the percentage change of the engagements w.r.t. a unit change in the independent variable. Table 2.13 summarizes the total effects of the independent variables on actual engagements, which is first computed for each endorser-task dyad and then averaged over the entire population. The relative partial effects of independent variables on potential outcome are taken directly from Table 2.11, as we have shown in Section 2.4.2 that  $\frac{\partial \log E[\mathbf{y}_{it}^* | \mathbf{x}_{it}]}{\partial x_{it}} = \boldsymbol{\beta}$ . The partial effects of independent variables on participation in Table 2.11 are also included in Table 2.13 to ease comparison.

For the majority of predictors, the direction of the total effects is consistent with that in the participation equation. Hence, participation is oftentimes the primary driver of actual engagements. Variables for which this holds include the number of followers, the verified ratio, the task number, and the number of days since registration. As a concrete example, though having participated in many campaigns is associated with low potential in generating engagements, such endorsers tend to generate an above-average number of actual engagements, due to their high tendency to participate.

Table 2.13 (Relative) Partial Effects on Participation, Potential and Actual Effectiveness

	Participation (Selection)			Potential Engagements (Outcome)			Actual Engagements (Overall)		
	likes	comments	retweets	likes	comments	retweets	likes	comments	retweets
<b>Exclusion Variable</b>									
isEligible	2.724***	2.723***	2.720***				6.229***	6.247***	6.198***
<b>Incentive</b>									
log(gain)	0.007	0.008	0.009	-0.039	-0.011	-0.061	-0.022	0.008	-0.041
log(loss)	-0.073***	-0.074***	-0.073***	0.022	0.146	0.095	-0.145	-0.024	-0.071
log(avgRwd)	0.034	0.036	0.037	-0.092	0.174	0.333	-0.014	0.256	0.417.
<b>Social Media Fanbase</b>									
log(followers)	0.083***	0.077***	0.079***	0.221.	-0.039	-0.195	0.411**	0.139	-0.015
verifiedRatio	0.256*	0.297**	0.283**	0.652	1.105	0.989	1.237.	1.788*	1.634*
<b>Prior Activity Level</b>									
log(tweetNum)	0.035.	0.045*	0.039.	0.001	-0.478***	-0.252*	0.082	-0.376**	-0.164
log(taskNum)	0.555***	0.564***	0.561***	-0.379***	-0.483***	-0.57***	0.889***	0.812***	0.708***
approvalRate	0.063	0.051	0.063	0.704	1.708.	-0.300	0.848	1.826.	-0.158
<b>Community Embeddedness</b>									
regDays	-4.375***	-4.41***	-4.358***	0.102	2.826*	3.258**	-9.902***	-7.293***	-6.673***
log(friends)	-0.088*	-0.084*	-0.086*	0.42**	0.131	0.368	0.219	-0.063	0.172
<b>Others</b>									
log(referralRwd)	-0.015.	-0.017.	-0.017.	0.023	0.003	-0.095	-0.013	-0.035	-0.134.
times	-0.142**	-0.142**	-0.147**	-0.285	-0.102	0.017	-0.611	-0.428	-0.317

However, participation doesn't always dominate potential. For example, the direction of the total effect of the number of tweets is more consistent with its direction in the outcome equation, rather than the selection equation. The effect size of tweet number in the selection equation is small compared to that in the outcome equation. This finding is not surprising once one realizes that the mean value of  $c_{it}$  in Equation (2.8) is 2.3 in our dataset. In some cases, the opposite effects in the participation and outcome equations may cancel out in the total effects, such as the total effect of the number of weuitui friends. These findings suggest that neglecting either participation or potential effectiveness in marketing campaigns can result in wrong decisions to target particular kinds of endorsers.

Advertisers unable to increase the participation or effectiveness of given endorsers may want to target endorsers who are likely to generate actual engagements, such as those who have participated in many tasks previously, have large fan bases, or a high verified ratio.

## 2.8 Conclusions

Paid endorsement, as an affordable approach to social advertising, has gained popularity among small firms in recent years. However, little is known on how to effectively target and incent paid endorsers. This paper provides new insights on how incentives and endorser characteristics affect participation and effectiveness. We conduct a field experiment on one of the largest paid endorsement platforms in China. For identification purpose, we exogenously manipulated the incentive and eligibility for participation. In order to analyze the collected panel count data on customer engagements while accounting for self-selection and repeated observations, we propose an approach that can address both challenges simultaneously.

Four findings have important implications for paid endorsement campaigns. First, endorsers are sensitive to losses but not gains as compared to their average reward in the past. This means that providing financial incentives in excess to an endorser's average reward per task in the past is very likely a waste of money. Advertisers should give attention to other aspects of the campaigns, such as non-monetary motives and the content of the ad message.

Second, the propensity to participate and effectiveness in generating engagements are often at odds with each other. This is so for both observed and unobserved characteristics. Consequently, it is difficult to find endorsers who are both responsive and effective. Advertisers should explore ways identify eligibility requirements that attract endorsers who are effective or find ways to boost the effort and effectiveness of endorsers who are responsive but otherwise ineffective. This may involve offering tasks with different requirements, eligibility restrictions or ad messages to different types of endorsers.

Third, it is misleading to assess the quality of endorsers solely based on the observed (actual) engagements. Endorsers observed to generate high engagements are not necessarily the most effective, but may simply be the most likely to participate. Conversely, the most effective endorsers tend not to participate in campaigns very often. This type of "latent gold" endorsers may easily be

overlooked by marketers who do not distinguish between participation and effectiveness, or between actual and potential effectiveness.

Finally, which endorsers to target depends on the objectives of marketers. Some endorser characteristics are associated with generating higher-effort engagements such as comments and retweets, whereas other are associated with lower-effort engagements such as likes.

Our work opens up several interesting directions for future research. First, it would be useful to study the effectiveness of endorsers in generating sales. Unfortunately, sales are hard to track at the individual level. Currently, the standard way to track clicks and sales is to use different short URLs in different tweets (even if they are for the same product), such that the source of clicks and sales can be tracked back to the short URLs. This means that clicks and sales can only be monitored at the level of task rather than endorser. More fine-grained tracking techniques are needed to study the effectiveness of individual endorsers in generating clicks and sales.

Secondly, our findings suggest that it would be useful to investigate the cost-effectiveness of different types of endorsers. This may require varying incentives or pay rates over a broad range to get robust insights. It may also be interesting to study how the effectiveness of various incentive and targeting approaches vary across product categories that vary in the status enhancement they provide to endorsers, such as mass vs. niche products or utilitarian vs. hedonic products.

Finally, the composition of the original message posted by the advertiser may also be worth investigating, as effective copy needs to appeal both to endorsers and to their followers. Here again, the distinction between participation and effectiveness may be essential to generating new fine-grained insights.

# APPENDIX

## Appendix 1.1: Simulation Studies

Our proposed model is called a collective cause model because it rests on the assumption that the event is caused by all co-senders collectively. We test the performance of the model in recovering the true parameters when the data are generated under the collective cause assumption. In practice, it is also possible that only part of the co-senders contributes to the event. To demonstrate the effectiveness of the collective cause model in dealing with such data, we focus on an extreme case in which the event is caused by one of the co-senders independently (called single-cause data). For simplicity, we assume that all the co-senders of a receiver adopt simultaneously at the beginning. This assumption has no effect on identification but greatly simplifies the data generation process. To test the robustness of the collective cause model to the distribution of independent variables, we assume that every user has three attributes drawn from three different distributions, namely, normal, binomial, and exponential. With a goal to generate a dataset with 10K events, we construct the collective-cause and single-cause datasets as follows:

- 1) Generate 200 senders and 5000 receivers, each has three attributes drawn from three different distributions: one normal, one binomial, and one exponential.
- 2) Randomly sample 10,000 senders and 10,000 receivers with replacement from the pool of 200 senders and 5,000 receivers, respectively. A one-to-one mapping between the 10K senders and 10K receivers results in 10K dyadic observations.
- 3) Randomly sample another 2,000 senders with replacement from the pool of 200 senders and map each of them to one of the 10K receivers in step (2) randomly. Those matched receivers in this step will therefore have multiple senders.
- 4) For each dyadic observation, compute the dyadic hazard, assuming the baseline hazard and all model parameters equal to 1.

5) *Collective-cause*: for each of the 10K receivers, compute her aggregated hazards by summing up the hazards from all her co-senders. Simulate a survival time for each receiver based on her aggregate hazards (Bender et al. 2005).

*Single-cause*: simulate a survival time for each of the 12K dyadic observations, following the method proposed by Bender et al. (Bender et al. 2005). If a receiver has multiple survival times associated with multiple senders, choose the minimum survival time as the survival time of the receiver.

6) To make the data more realistic, choose the lower 20% quantile of all survival times as the censoring time, such that 80% of conversion events are censored in the final data.

The data generation process of the collective-cause and single-cause data are exactly the same, except for step (5). We use a dyadic setup to ensure that the structure of the simulated dataset is similar to the structure of the dataset used in the application. Moreover, we censored 80% of events to test the effectiveness of the single cause model on incomplete observations.

To show the effectiveness of the proposed collective cause model, which doesn't speculate on the quantitative contribution of co-senders, we compared its performance with two benchmark models developed based on the idea of linear attribution in advertising.<sup>30</sup> The key idea of linear attribution is that each touch point contributes equally to the conversion. In the first benchmark model, we assume that every co-sender has equal probability to be the sole cause of event and maximize the expected likelihood of the event to be caused by any co-sender. We call this model the equal probability model. In contrast to the first benchmark model which assumes that only one of the co-senders is the true cause, in the second model we assume that every co-sender is part of the true cause. Specifically, we treat an event with multiple co-senders as multiple independent events caused by the co-senders each. We restrict the total case weight of each receiver to be one and evenly split the unit case weight among multiple co-senders. The second benchmark model is

---

<sup>30</sup> <https://support.google.com/analytics/answer/1662518?hl=en>

called the tied events model as it can be estimated by the tie handling methods of proportional hazard models (Therneau 2000).

Table A1.1 summarizes the relative errors (i.e.,  $\frac{\hat{\beta} - \beta}{\beta}$ ) of three models on two types of datasets, averaged over 20 runs. The prefix “r” indicates covariates on the receiver side. Enclosed in parentheses are the standard deviations of the relative errors.

Table A1.1 Relative Errors of the Collective Cause Model

	Single-Cause Data			Collective-Cause Data		
	Tied Events	Equal Prob.	Collective Cause	Tied Events	Equal Prob.	Collective Cause
normal	-0.1883 (0.02)	-0.2102 (0.03)	0.0064 (0.02)	-0.3312 (0.02)	-0.2122 (0.03)	0.0014 (0.02)
binomial	-0.1787 (0.04)	-0.1996 (0.05)	0.0060 (0.05)	-0.3236 (0.04)	-0.2079 (0.04)	-0.0053 (0.04)
exponential	-0.1539 (0.01)	-0.1776 (0.01)	-0.0006 (0.02)	-0.2729 (0.02)	-0.1766 (0.01)	-0.0003 (0.02)
rnormal	-0.1897 (0.02)	-0.2109 (0.02)	0.0030 (0.02)	-0.3254 (0.02)	-0.2079 (0.02)	0.0066 (0.02)
rbinomial	-0.1811 (0.05)	-0.2045 (0.05)	-0.0034 (0.05)	-0.3200 (0.05)	-0.2088 (0.05)	-0.0064 (0.04)
rexponential	-0.1541 (0.01)	-0.1758 (0.01)	0.0015 (0.01)	-0.2740 (0.01)	-0.1743 (0.01)	0.0022 (0.02)

As can be seen, the proposed collective cause model can recover the true parameters with negligible errors not only on the collective-cause data, but also on the single-cause data. This finding demonstrates that the collective cause model is a valid model even if only part of the co-senders contributes to the event. The mathematical proof regarding why the collective cause model can still recover the true parameters when only one of co-senders contributes to the event is available from the authors upon request. The intuition behind this finding is that, in the single-cause data, the overall hazard of a receiver given in the numerator of Equation (2.2) can be reinterpreted as the overall hazard of the receiver to be influenced by any single source she has seen. In this sense, the collective cause model is a truthful representation of the single cause data, except that it does not use the true cause information. The estimates of the tied events model and equal probability model are both substantially biased downwards, which demonstrates that arbitrary assignment of credits among co-senders may lead to misleading results. The effectiveness of the collective cause model in recovering the true parameters are robust to censoring, scaling, distribution of survival times, and average number of co-senders on a receiver.

## Appendix 1.2: Complete Results

Table A1.2 Complete Parameter Estimates for Models in Table 1.8

	Model1	Model2	Model3	Model4
<b>Characteristics of Sender</b>				
isSocialTRUE	1.409**	1.577***	1.602***	1.482**
isDiggAdsTRUE	0.36	-0.434	0.277	0.005
logFollowees	-0.003	0.016	-0.015	-0.004
logFollowers	-0.872***	-0.759***	-0.827***	-0.798***
logMutuals	0.002	0.011	0.026	0.028
logDiggs	-0.32*	-0.285*	-0.269.	-0.239
logComments	-0.191.	-0.037	-0.183.	-0.155
logSubmissions	-0.009	-0.148	-0.005	-0.069
logAvgDiggs	0.585***	0.445***	0.444***	0.417***
logAvgComments	0.008	-0.105	0.025	-0.001
logAvgSubmissions	0.04	0.121	0.038	0.081
genderf	0.222	-0.047	-0.096	-0.089
genderm	0.274*	0.274*	0.209.	0.231.
<b>Characteristics of Receiver</b>				
logFollowees	-0.238***	-0.248***	-0.246***	-0.248***
logFollowers	-0.208***	-0.206***	-0.211***	-0.209***
logMutuals	-0.101***	-0.109***	-0.102***	-0.104***
logDiggs	0.166***	0.169***	0.171***	0.17***
logComments	-0.156***	-0.16***	-0.163***	-0.162***
logSubmissions	-0.152***	-0.153***	-0.155***	-0.154***
logAvgDiggs	0.438***	0.435***	0.435***	0.435***
logAvgComments	0.218***	0.217***	0.224***	0.221***
logAvgSubmissions	0.029	0.031	0.031	0.03
genderf	0.12**	0.129**	0.129**	0.13**
genderm	0.137***	0.143***	0.143***	0.144***
<b>Sharing Timing of Sender</b>				
wday1	-0.462**	-0.324*	-0.448**	-0.382*
wday2	0.119	-0.117	-0.11	-0.109
wday3	0.263*	0.073	0.122	0.088
wday4	0.129	0.048	0.045	0.039
wday5	-0.103	0.03	-0.002	0.021
wday6	0.117	0.265	0.004	0.148
hour(5,11]	-0.348**	-0.281*	-0.259*	-0.268*
hour(11,17]	-0.256*	-0.165	-0.196	-0.179
hour(17,23]	0.044	0.006	0.003	-0.002
shareTime	-0.023	0.099**	0.097**	0.116**
<b>Number of Co-senders</b>				
co-senders	-0.082***	-0.057***	-0.058***	-0.056***
<b>Dyadic Characteristics</b>				
isMutualTrue	-0.645***	-0.499***	-0.526***	-0.5***
logCommonFollowees	0.23***	0.174**	0.175**	0.175**
logCommonFollowers	0.845***	1.364***	0.829***	1.074***
logCommonMutuals	-0.245***	-0.19***	0.799***	0.418**
logCommonFollowers:logPopularity		-0.153***		-0.071**
logCommonMutuals:logPopularity			-0.258***	-0.16***
<b>Fitness</b>				
logLikelihood	-22661	-22623	-22620	-22618
AIC	45401	45326	45320	45317

The three levels for gender are: m – male, f – female, and u – unknown. For wday, Monday is coded as 0. Hour of a day is grouped into four bins. For dummy variables, the missing levels are the reference levels.

## Appendix 1.3 Diffusion Graphs

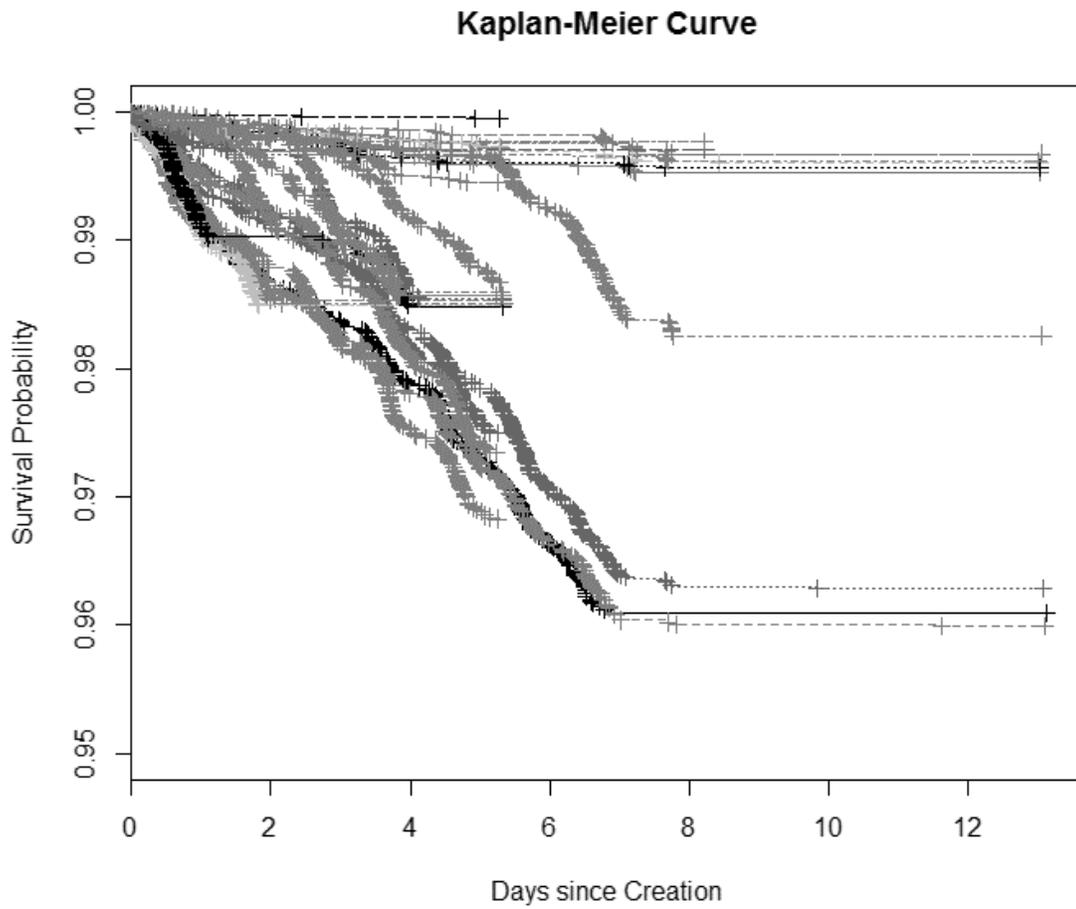


Figure A1.1 Kaplan-Meier Survival Curve for Digg Ads<sup>31</sup>

---

<sup>31</sup> The KM curve is computed based on the average survival probability of all receivers who are at risk over time

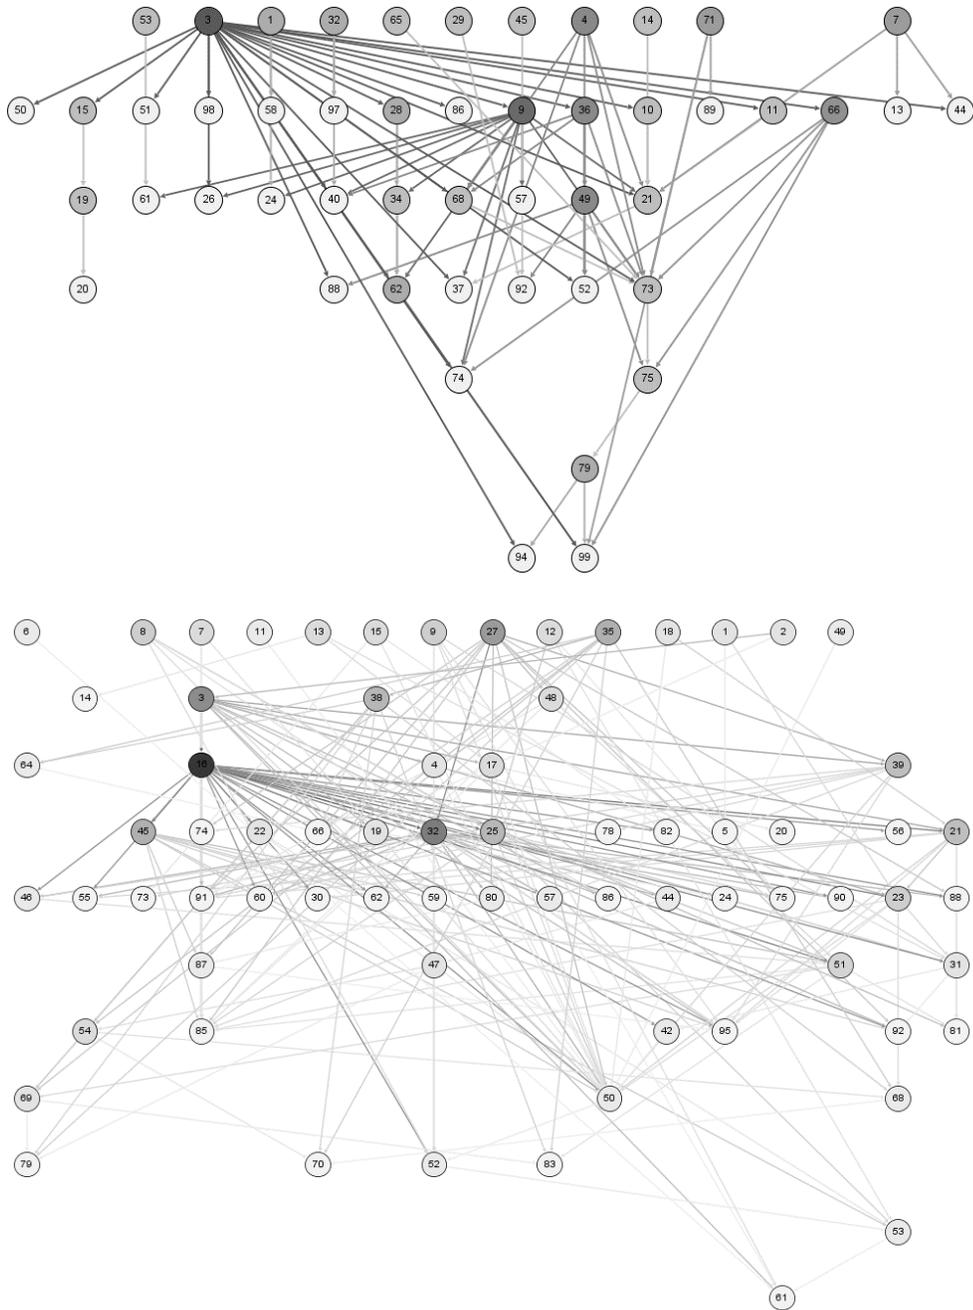


Figure A1.2 Sharing Graphs for Ads 1 & 2.<sup>32</sup>

<sup>32</sup> Arrow represents information flow. Nodes without incoming links share spontaneously. Nodes are labeled based on the order they share the ad. The darkness of the color of a node is proportional to her outgoing links in the graph. The color of an arrow is consistent with the source node.

## Appendix 1.4 Additional Results

Table A1.3 Parameter Estimates on the Digg Dataset (including inactive users)

	Model1	Model2	Model3	Model4
<b>Embeddedness</b>				
logCommonFollowees	0.257***	0.200***	0.208***	0.205***
logCommonFollowers	0.723***	1.275***	0.710***	1.043***
logCommonMutuals	-0.227***	-0.179***	0.754***	0.290**
<b>Interactions with Popularity</b>				
logCommonFollowers:logPopularity		-0.152***		-0.091***
logCommonMutuals:logPopularity			-0.237***	-0.121***
<b>Fitness</b>				
logLikelihood	-23887	-23843	-23845	-23839
AIC	47844	47757	47762	47752

Table A1.4 Summary Statistics for Twitter Dataset

Number of ads/tweets	74
Number of sharing users (senders)	4,209
Number of potential sharing users (receivers)	36,187
Number of <sender, receiver> dyads	90,288
Number of <sender, receiver, ad> tuples	171,685
Number of spontaneous tuples	80,721 (47%)
Number of social tuples	90,964 (53%)
Number of shares (retweets)	4,740
Number of spontaneous shares	1,020 (21.5%)
Number of potential influenced shares	3,720 (78.5%)
Percentage with more than one co-senders (excluding special sender)	6.8%

Table A1.5 Descriptions of Independent Variables for Twitter Dataset

Independent Variable		Description
$X_i/X_j$		<b>Attributes of sender <math>i</math> / receiver <math>j</math></b>
Network attributes	followees	Number of followees (out-degree)
	followers	Number of followers (in-degree)
	mutuals	Number of mutual followers
	lists	Number of lists subscribed
Engagement levels	statuses	Total number of tweets, including retweets
	favourites	Total number of favourites
Others	verified	Whether the Twitter account is verified
	regMon	How many months have the user been registered on Twitter
	isSocial ( $s_i$ )	1 if sender $i$ is a social source (i.e., followee), otherwise 0
	isAuthor	1 if the sender is the author of the tweet, otherwise 0
$X_{ij}$		<b>Attributes of a sender-receiver dyad</b>
Dyadic network attributes	isMutual	Does the sender and the receiver follow each other mutually
	commonFollowees	Number of followees shared by the sender and the receiver
	commonFollowers	Number of followers shared by the sender and the receiver
	commonMutuals	Number of mutual followers shared by the sender and the receiver
$X_{ik}$		<b>Sender-specific attributes of a tweet</b>
Sharing timing	wday	Day of a week when sender $i$ retweeted tweet $k$
	hour	Hour of a day when sender $i$ retweeted tweet $k$
	shareTime	Hours taken for sender $i$ to retweet since creation of tweet $k$ , 0 for the front page
$X_{jk}$		<b>Receiver-specific attributes of a tweet</b>
	co-senders	Number of followees (co-senders) of the receiver who have already shared
$X_k$		<b>Attributes of ads <math>k</math> (only interaction with other variables can be identified)</b>
	popularity	Number of retweets at a given time point

Table A1.6 Key Statistics of Main Variables for Twitter Dataset

	Zeros	Mean	SD	Min	Median	Max
<b>Unitary Network Attributes of All Users</b>						
Number of followees	14	9298.8	43743.3	0	769	2422154
Number of followers	202	18843.9	348931.6	0	380	59159316
Number of mutuals	945	6278.0	31334.0	0	212	1755611
<b>Dyadic Network Attributes of Sender-Receiver Dyads</b>						
isMutual (1 – reciprocal, 0 – non-reciprocal)	40940	0.45	0.50	0	0	1
Number of common followees	9596	61.1	715.5	0	10	79376
Number of common followers	16015	235.3	6016.5	0	5	500406
Number of common mutual followers	26864	74.5	490.5	0	2	35478
<b>Popularity of Tweets</b>						
Number of retweets	15	20.5	41.9	0	10	379

Table A1.7 Correlation among Dyadic Network Characteristics for Twitter Dataset

	isMutual	logCommonFollowees	logCommonFollowers	logCommonMutuals
isMutual	1.00	0.10	0.12	0.50
logCommonFollowees		1.00	0.56	0.53
logCommonFollowers			1.00	0.63
logCommonMutuals				1.00

Table A1.8 Complete Results on the Twitter Dataset

	Model1	Model2	Model3	Model4
<b>Characteristics of Sender</b>				
isSocialTRUE	-1.575	-0.762	-0.742	-0.633
isAuthorTRUE	1.459***	1.069***	1.23***	1.076***
logFollowees	-0.764***	-0.938***	-0.97***	-0.997***
logFollowers	0.399***	0.416***	0.412***	0.421***
logMutuals	0.811***	0.99***	0.995***	1.042***
logLists	0.216*	0.161*	0.176**	0.154*
logStatuses	-0.142**	-0.147***	-0.119**	-0.136***
logFavourites	-0.478***	-0.463***	-0.467***	-0.463***
verified	-1.392*	-1.449***	-1.771***	-1.604***
logRegMon	0.482***	0.463***	0.416***	0.432***
<b>Characteristics of Receiver</b>				
logFollowees	-0.468***	-0.482***	-0.473***	-0.48***
logFollowers	-0.139***	-0.145***	-0.143***	-0.146***
logMutuals	-0.049**	-0.038*	-0.047**	-0.039*
logLists	0.099***	0.086***	0.091***	0.086***
logStatuses	0.278***	0.284***	0.281***	0.283***
logFavourites	0.081***	0.084***	0.084***	0.085***
verified	-0.525*	-0.702**	-0.685**	-0.772***
logRegMon	-0.165***	-0.163***	-0.164***	-0.163***
<b>Sharing Timing of Sender</b>				
wday0	-1.233	-0.976	-0.92	-0.936
wday1	-0.049	-0.147	-0.177	-0.186
wday2	-0.616	-0.5	-0.538	-0.491
wday3	2.106**	1.782**	1.753**	1.721**
wday5	1.387*	1.266**	1.468***	1.342**
wday6	0.23	-0.575	-0.568	-0.579
hour(5.75,11.5]	-2.347***	-2.449***	-2.479***	-2.513***
hour(11.5,17.2]	-1.829***	-1.476***	-1.559***	-1.458***
hour(17.2,23]	-0.253	-0.274	-0.223	-0.249
shareTime	0.054***	0.05***	0.051***	0.05***
<b>Number of Co-senders</b>				
co-senders	-1.865***	-2.072***	-1.993***	-2.077***
<b>Dyadic Characteristics</b>				
isMutualTRUE	0.162*	0.146.	0.164*	0.149*
logCommonFollowees	0.294***	0.311***	0.299***	0.309***
logCommonFollowers	-0.127***	0.227***	-0.132***	0.144***
logCommonMutuals	-0.035	0.013	0.643***	0.284***
logCommonFollowers:logPopularity		-0.105***		-0.081***
logCommonMutuals:logPopularity			-0.174***	-0.075***
<b>Fitness</b>				
likelihood	-29378	-29324	-29340	-29319
AIC	58822	58717	58747	58708

For wday, Monday is coded as 0. Hour of a day is grouped into four bins. For dummy variables, the missing levels are the reference levels.

## Appendix 2.1 Likelihood and Parameter Estimation

The likelihood of all observations on endorser  $i$  is given by

$$L_i = P(y_{i1}^*, \dots, y_{iT}^*; z_{i1}, \dots, z_{iT} | \mathbf{x}_{i1}, \dots, \mathbf{x}_{iT}, \mathbf{w}_{i1}, \dots, \mathbf{w}_{iT})$$

$$= \int_{-\infty}^{\infty} \phi(\varepsilon_i) d\varepsilon_i \int_{-\infty}^{\infty} f(u_i | \varepsilon_i) du_i \int_{-\infty}^{\infty} P(y_{i1}^*, \dots, y_{iT}^*; z_{i1}, \dots, z_{iT} | \mathbf{x}_{i1}, \dots, \mathbf{x}_{iT}, \mathbf{w}_{i1}, \dots, \mathbf{w}_{iT}, \varepsilon_i, u_i, \varepsilon_{i1}, \dots, \varepsilon_{iT}) \phi(\varepsilon_{i1}) \dots \phi(\varepsilon_{iT}) d\varepsilon_{i1} \dots d\varepsilon_{iT}$$

Based on the i.i.d. assumption on the error terms  $\varepsilon_{i1}, \dots, \varepsilon_{iT}$ , the above likelihood can be simplified as

$$L_i = \int_{-\infty}^{\infty} \phi(\varepsilon_i) d\varepsilon_i \int_{-\infty}^{\infty} f(u_i | \varepsilon_i) du_i \prod_{t=1}^{T_i} \int_{-\infty}^{\infty} P(y_{it}^* | \mathbf{x}_{it}, \varepsilon_i, \varepsilon_{it})^{z_{it}} P(z_{it} | \mathbf{w}_{it}, u_i, \varepsilon_{it}) \phi(\varepsilon_{it}) d\varepsilon_{it}$$

$$= \frac{1}{2\pi\sqrt{1-\rho^2}} \int_{-\infty}^{\infty} e^{-\frac{\varepsilon_i^2}{2}} d\varepsilon_i \int_{-\infty}^{\infty} e^{-\frac{(u_i - \rho\varepsilon_i)^2}{2(1-\rho^2)}} du_i \prod_{t=1}^{T_i} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} P(y_{it}^* | \mathbf{x}_{it}, \varepsilon_i, \varepsilon_{it})^{z_{it}} \Phi\left((2z_{it} - 1) \frac{\alpha\mathbf{w}'_{it} + \delta u_i + \tau\varepsilon_{it}}{\sqrt{1-\tau^2}}\right) e^{-\frac{\varepsilon_{it}^2}{2}} d\varepsilon_{it}$$

After substituting in  $r = \frac{\varepsilon_i}{\sqrt{2}}$ ,  $s = \frac{u_i - \rho\varepsilon_i}{\sqrt{2(1-\rho^2)}}$ , and  $v = \frac{\varepsilon_{it}}{\sqrt{2}}$ , the likelihood can be further simplified as

$$L_i = \frac{1}{\pi} \int_{-\infty}^{\infty} e^{-r^2} dr \int_{-\infty}^{\infty} e^{-s^2} ds \prod_{t=1}^{T_i} \frac{1}{\sqrt{\pi}} \int_{-\infty}^{\infty} P(y_{it}^* | \mathbf{x}_{it}, \sqrt{2}r, \sqrt{2}v)^{z_{it}} \Phi\left((2z_{it} - 1) \frac{\alpha\mathbf{w}'_{it} + \sqrt{2}\rho\delta r + \sqrt{2(1-\rho^2)}\delta s + \sqrt{2}\tau v}{\sqrt{1-\tau^2}}\right) e^{-v^2} dv$$

In this form, the likelihood can be approximated numerically by Gauss-Hermite quadrature (Abramowitz and Stegun 1972; Greene 2009). Using the Gauss-Hermite quadrature procedure three times, the likelihood can be approximated as

$$L_i \approx \frac{1}{\pi} \sum_{h=1}^H \omega_h \sum_{k=1}^K \varphi_k \prod_{t=1}^{T_i} \sum_{m=1}^M \frac{\mu_m}{\sqrt{\pi}} P(y_{it}^* | \mathbf{x}_{it}, \sqrt{2}r, \sqrt{2}v)^{z_{it}} \Phi\left((2z_{it} - 1) \frac{\alpha\mathbf{w}'_{it} + \sqrt{2}\rho\delta r_h + \sqrt{2(1-\rho^2)}\delta s_k + \sqrt{2}\tau v_m}{\sqrt{1-\tau^2}}\right)$$

where  $\{w_h, r_h\}$ ,  $\{\varphi_k, s_k\}$ ,  $\{\mu_m, v_m\}$  represent the weights and nodes for the three quadratures, with  $H$ ,  $K$ , and  $M$  points being used respectively. In our analysis, we find that  $H = K = M = 10$  are sufficient to yield reasonably good approximations.

The overall log likelihood on the entire dataset can be written as

$$LL = \ln \prod_{i=1}^N L_i = \sum_{i=1}^N \ln L_i$$

Since  $\rho, \tau \in [-1, 1]$ , we use the L-BFGS-B method (Byrd et al. 1995; Zhu et al. 1997), which allows for box constraints, to maximize the log likelihood. We have found from simulations that the above likelihood maximization procedure can recover the true parameters very well.

Due to the complexity of the log likelihood, the asymptotic covariance matrix of the parameter estimates is computed using the BHHH estimator (Berndt et al. 1974; Greene 2009), which only requires the computation of the score function (i.e., gradient of log likelihood).

Let  $\theta = [\alpha, \delta, \rho, \tau]$  and  $\eta = [\beta, \sigma, \gamma]$ , the gradient of the log likelihood is approximated by

$$\frac{\partial LL}{\partial \theta} \approx \frac{1}{\pi} \sum_{i=1}^N \frac{1}{L_i} \sum_{h=1}^H \omega_h \sum_{k=1}^K \varphi_k \left( \prod_{t=1}^{T_i} \sum_{m=1}^M \frac{\mu_m}{\sqrt{\pi}} P(y_{it} | \lambda_{it}^{hm})^{z_{it}} \Phi(q_{it}^{hkm}) \right) \sum_{t=1}^{T_i} \frac{\sum_{m=1}^M \frac{\mu_m}{\sqrt{\pi}} P(y_{it} | \lambda_{it}^{hm})^{z_{it}} \Phi(q_{it}^{hkm})^{z_{it}-1} w_{it}^{ext}}{\sum_{m=1}^M \frac{\mu_m}{\sqrt{\pi}} P(y_{it} | \lambda_{it}^{hm})^{z_{it}} \Phi(q_{it}^{hkm})}$$

$$\text{where } w_{it}^{ext} = \left[ w_{it}, \sqrt{2\rho r_h} + \sqrt{2(1-\rho^2)} s_k, \sqrt{2} \delta r_h - \sqrt{\frac{2}{(1-\rho^2)}} \rho \delta s_k, \frac{\sqrt{2} v_m + (\alpha w'_{it} + \sqrt{2\rho} \delta r_h + \sqrt{2(1-\rho^2)} \delta s_k) \tau}{1-\tau^2} \right]$$

$$\frac{\partial LL}{\partial \eta} \approx$$

$$\frac{1}{\pi} \sum_{i=1}^N \frac{1}{L_i} \sum_{h=1}^H \omega_h \sum_{k=1}^K \varphi_k \left( \prod_{t=1}^{T_i} \sum_{m=1}^M \frac{\mu_m}{\sqrt{\pi}} P(y_{it} | \lambda_{it}^{hm})^{z_{it}} \Phi(q_{it}^{hkm}) \right) \sum_{t=1}^{T_i} \frac{\sum_{m=1}^M z_{it} \frac{\mu_m}{\sqrt{\pi}} P(y_{it} | \lambda_{it}^{hm}) (y_{it} - \lambda_{it}^{hm}) \Phi(q_{it}^{hkm}) x_{it}^{ext}}{\sum_{m=1}^M \frac{\mu_m}{\sqrt{\pi}} P(y_{it} | \lambda_{it}^{hm})^{z_{it}} \Phi(q_{it}^{hkm})}$$

$$\text{where } x_{it}^{ext} = [x_{it}, \sqrt{2} r_h, \sqrt{2} v_m].$$

## Appendix 2.2 Mean Actual Outcome

The unconditional (not conditional on any unobserved variable value) mean of actual outcome can be derived as follows

$$\begin{aligned}
E[y_{it}|\mathbf{x}_{it}, \mathbf{w}_{it}] &= E_{u_i} E_{\varepsilon_i|u_i} E_{\varepsilon_{it}} [E[y_{it}|\mathbf{x}_{it}, \mathbf{w}_{it}, u_i, \varepsilon_i, \varepsilon_{it}]] \\
&= E_{u_i} E_{\varepsilon_i|u_i} E_{\varepsilon_{it}} [P(z_{it} = 1|\mathbf{w}_{it}, u_i, \varepsilon_{it}) E[y_{it}^*|\mathbf{x}_{it}, \varepsilon_i, \varepsilon_{it}]] \\
&= \int_{-\infty}^{\infty} \phi(u_i) du_i \int_{-\infty}^{\infty} f(\varepsilon_i|u_i) d\varepsilon_i \int_{-\infty}^{\infty} \Phi\left(\frac{\alpha\mathbf{w}'_{it} + \delta u_i + \tau\varepsilon_{it}}{\sqrt{1-\tau^2}}\right) \exp(\beta\mathbf{x}'_{it} + \sigma\varepsilon_i + \gamma\varepsilon_{it}) \phi(\varepsilon_{it}) d\varepsilon_{it} \\
&= \exp(\beta\mathbf{x}'_{it}) \int_{-\infty}^{\infty} \left[ \int_{-\infty}^{\infty} \Phi\left(\frac{\alpha\mathbf{w}'_{it} + \delta u_i + \tau\varepsilon_{it}}{\sqrt{1-\tau^2}}\right) \exp(\gamma\varepsilon_{it}) \phi(\varepsilon_{it}) d\varepsilon_{it} \right] \phi(u_i) du_i \int_{-\infty}^{\infty} \exp(\sigma\varepsilon_i) f(\varepsilon_i|u_i) d\varepsilon_i \\
&= \exp(\beta\mathbf{x}'_{it}) \int_{-\infty}^{\infty} \left[ \int_{-\infty}^{\infty} \Phi\left(\frac{\alpha\mathbf{w}'_{it} + \delta u_i + \tau\varepsilon_{it}}{\sqrt{1-\tau^2}}\right) \exp(\gamma\varepsilon_{it}) \phi(\varepsilon_{it}) d\varepsilon_{it} \right] \phi(u_i) du_i \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi(1-\rho^2)}} e^{\rho\sigma u_i + \frac{(1-\rho^2)\sigma^2}{2}} e^{-\frac{(\varepsilon_i - \rho u_i - (1-\rho^2)\sigma)^2}{2(1-\rho^2)}} d\varepsilon_i \\
&= \exp(\beta\mathbf{x}'_{it}) \int_{-\infty}^{\infty} \left[ \int_{-\infty}^{\infty} \Phi\left(\frac{\alpha\mathbf{w}'_{it} + \delta u_i + \tau\varepsilon_{it}}{\sqrt{1-\tau^2}}\right) \exp(\gamma\varepsilon_{it}) \phi(\varepsilon_{it}) d\varepsilon_{it} \right] e^{\rho\sigma u_i + \frac{(1-\rho^2)\sigma^2}{2}} \phi(u_i) du_i \\
&= \frac{1}{2\pi} \exp\left(\beta\mathbf{x}'_{it} + \frac{\sigma^2 + \gamma^2}{2}\right) \int_{-\infty}^{\infty} \left[ \int_{-\infty}^{\infty} \Phi\left(\frac{\alpha\mathbf{w}'_{it} + \delta u_i + \tau\varepsilon_{it}}{\sqrt{1-\tau^2}}\right) e^{-\frac{(\varepsilon_{it} - \gamma)^2}{2}} d\varepsilon_{it} \right] e^{-\frac{(u_i - \rho\sigma)^2}{2}} du_i
\end{aligned}$$

Let  $v = \frac{\varepsilon_{it} - \gamma}{\sqrt{2}}$  and  $r = \frac{u_i - \rho\sigma}{\sqrt{2}}$ , the above equation can be simplified as

$$E[y_{it}|\mathbf{x}_{it}, \mathbf{w}_{it}] = \frac{1}{\pi} \exp\left(\beta\mathbf{x}'_{it} + \frac{\sigma^2 + \gamma^2}{2}\right) \int_{-\infty}^{\infty} \left[ \int_{-\infty}^{\infty} \Phi\left(\frac{\alpha\mathbf{w}'_{it} + \sqrt{2}\delta r + \rho\sigma\delta + \tau\gamma + \sqrt{2}\tau v}{\sqrt{1-\tau^2}}\right) e^{-v^2} dv \right] e^{-r^2} dr$$

Similar to the likelihood function,  $E[y_{it}|\mathbf{x}_{it}, \mathbf{w}_{it}]$  can be approximated by two embedded Gauss-Hermite quadratures.

$$E[y_{it}|\mathbf{x}_{it}, \mathbf{w}_{it}] \approx \frac{1}{\pi} \exp\left(\beta\mathbf{x}'_{it} + \frac{\sigma^2 + \gamma^2}{2}\right) \sum_{h=1}^H \omega_h \sum_{k=1}^K \varphi_k \Phi\left(\frac{\alpha\mathbf{w}'_{it} + \sqrt{2}\delta r_h + \rho\sigma\delta + \tau\gamma + \sqrt{2}\tau v_k}{\sqrt{1-\tau^2}}\right)$$

where  $\{w_h, r_h\}$ , and  $\{\varphi_k, s_k\}$  represent the weights and nodes for the two quadratures, with  $H$  and  $K$  points being used, respectively.

## Appendix 2.3 Relative Partial Effects

Letting  $q_{it}^{hk} = \frac{\alpha w'_{it} + \sqrt{2}\delta r_h + \rho\sigma\delta + \tau\gamma + \sqrt{2}\tau v_k}{\sqrt{1-\tau^2}}$ , the relative partial effects of  $x_{it}$  and  $w_{it}$  on the mean actual

outcome are given by

$$\mathbf{g}_{w_{it}} = \frac{\partial \log E[y_{it} | \mathbf{x}_{it}, \mathbf{w}_{it}]}{\partial \mathbf{w}_{it}} = c_{it} \boldsymbol{\alpha} \approx \frac{\sum_{h=1}^H \omega_h \sum_{k=1}^K \varphi_k \phi(q_{it}^{hk})}{\sqrt{1-\tau^2} \sum_{h=1}^H \omega_h \sum_{k=1}^K \varphi_k \Phi(q_{it}^{hk})} \boldsymbol{\alpha}$$

$$\mathbf{g}_{x_{it}} = \frac{\partial \log E[y_{it} | \mathbf{x}_{it}, \mathbf{w}_{it}]}{\partial \mathbf{x}_{it}} \approx \boldsymbol{\beta}$$

$$\text{where } c_{it} = \frac{\int_{-\infty}^{\infty} \left[ \int_{-\infty}^{\infty} \Phi\left(\frac{\alpha w'_{it} + \sqrt{2}\delta r + \rho\sigma\delta + \tau\gamma + \sqrt{2}\tau v}{\sqrt{1-\tau^2}}\right) e^{-v^2} dv \right] e^{-r^2} dr}{\sqrt{1-\tau^2} \int_{-\infty}^{\infty} \left[ \int_{-\infty}^{\infty} \Phi\left(\frac{\alpha w'_{it} + \sqrt{2}\delta r + \rho\sigma\delta + \tau\gamma + \sqrt{2}\tau v}{\sqrt{1-\tau^2}}\right) e^{-v^2} dv \right] e^{-r^2} dr} > 0.$$

If a variable appears in both  $\mathbf{x}_{it}$  and  $\mathbf{w}_{it}$ , the derivative for that variable would be the sum of the corresponding elements in  $\mathbf{g}_{x_{it}}$  and  $\mathbf{g}_{w_{it}}$ .

The covariance matrix of the relative partial effects can be approximated by the delta method.

Letting  $\boldsymbol{\theta} = [\boldsymbol{\alpha}, \boldsymbol{\beta}, \delta, \sigma, \gamma, \rho, \tau]$ ,  $A = \sum_{h=1}^H \omega_h \sum_{k=1}^K \varphi_k \phi(q_{it}^{hk})$ , and  $B = \sum_{h=1}^H \omega_h \sum_{k=1}^K \varphi_k \Phi(q_{it}^{hk})$ , the first order derivative of the relative partial effects w.r.t. the model parameters can be computed as

$$\begin{aligned} \mathbf{J}_{w_{it}} &= \frac{\partial \mathbf{g}'_{w_{it}}}{\partial \boldsymbol{\theta}} = \frac{A}{\sqrt{1-\tau^2} B} \frac{\partial \boldsymbol{\alpha}'}{\partial \boldsymbol{\theta}} + \boldsymbol{\alpha}' \frac{\partial \frac{A}{\sqrt{1-\tau^2} B}}{\partial \boldsymbol{\theta}} \\ &= \frac{A}{\sqrt{1-\tau^2} B} \left[ I_{|\boldsymbol{\alpha}'| \times |\boldsymbol{\alpha}|}, 0_{|\boldsymbol{\alpha}'| \times (|\boldsymbol{\beta}|+4)}, \frac{2\tau}{1-\tau^2} \boldsymbol{\alpha}' \right] + \boldsymbol{\alpha}' \frac{\sum_{h=1}^H \omega_h \sum_{k=1}^K -q_{it}^{hk} \varphi_k \phi(q_{it}^{hk}) \frac{\partial q_{it}^{hk}}{\partial \boldsymbol{\theta}} - \frac{A}{B} \sum_{h=1}^H \omega_h \sum_{k=1}^K \varphi_k \phi(q_{it}^{hk}) \frac{\partial q_{it}^{hk}}{\partial \boldsymbol{\theta}}}{\sqrt{1-\tau^2} B} \end{aligned}$$

where  $\frac{\partial q_{it}^{hk}}{\partial \boldsymbol{\theta}} = \frac{1}{\sqrt{1-\tau^2}} \left[ \mathbf{w}_{it}, 0_{|\boldsymbol{\beta}|}, \sqrt{2}r_h + \rho\sigma, \rho\delta, \tau, \sigma\delta, \frac{(\gamma + \sqrt{2}v_k)(1+\tau^2) + 2\tau(\alpha w'_{it} + \sqrt{2}\delta r_h + \rho\sigma\delta)}{1-\tau^2} \right]$

$$\mathbf{J}_{x_{it}} = \frac{\partial \mathbf{g}'_{x_{it}}}{\partial \boldsymbol{\theta}} = \frac{\partial \boldsymbol{\beta}'}{\partial \boldsymbol{\theta}} = [0_{|\boldsymbol{\beta}'| \times |\boldsymbol{\alpha}|}, I_{|\boldsymbol{\beta}'| \times |\boldsymbol{\beta}|}, 0_{|\boldsymbol{\beta}'| \times 5}]$$

Similarly, if a variable appears in both  $x_{it}$  and  $w_{it}$ , the derivative for that variable would be the sum of the corresponding rows in  $\mathbf{J}_{x_{it}}$  and  $\mathbf{J}_{w_{it}}$ .

Letting  $\mathbf{g}$  represent the relative partial effects of all variables, and  $\mathbf{J}$  be the first order derivative of  $\mathbf{g}$  w.r.t. the model parameters, then according to the delta method

$$\text{Var}(\mathbf{g}) \approx \mathbf{J}\mathbf{V}\mathbf{J}'$$

where  $\mathbf{V}$  is the covariance matrix of the parameter estimates.

## Appendix 2.4 Prediction of Participation and Effectiveness

1) Probability to participate

$$\begin{aligned}
 P(z_{it} = 1 | \mathbf{w}_{it}) &= \int_{-\infty}^{\infty} P(z_{it} = 1 | \mathbf{w}_{it}, u_i) \phi(u_i) du_i = \int_{-\infty}^{\infty} \Phi(\boldsymbol{\alpha} \mathbf{w}'_{it} + \delta u_i) \phi(u_i) du_i \\
 &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \Phi(\boldsymbol{\alpha} \mathbf{w}'_{it} + \delta u_i) e^{-\frac{u_i^2}{2}} du_i = \frac{1}{\sqrt{\pi}} \int_{-\infty}^{\infty} \Phi(\boldsymbol{\alpha} \mathbf{w}'_{it} + \sqrt{2}\delta v) e^{-v^2} dv = \frac{1}{\sqrt{\pi}} \sum_{k=1}^K \omega_k \Phi(\boldsymbol{\alpha} \mathbf{w}'_{it} + \sqrt{2}\delta v)
 \end{aligned}$$

2) Expected number of engagements

$$\begin{aligned}
 E[y_{it} | \mathbf{x}_{it}] &= E_{\varepsilon_i} E_{\varepsilon_{it}} [E[y_{it} | \mathbf{x}_{it}, \varepsilon_i, \varepsilon_{it}]] = \int_{-\infty}^{\infty} \phi(\varepsilon_i) d\varepsilon_i \int_{-\infty}^{\infty} \exp(\boldsymbol{\beta} \mathbf{x}'_{it} + \sigma \varepsilon_i + \gamma \varepsilon_{it}) \phi(\varepsilon_{it}) d\varepsilon_{it} \\
 &= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\frac{\varepsilon_i^2}{2}} d\varepsilon_i \int_{-\infty}^{\infty} \exp(\boldsymbol{\beta} \mathbf{x}'_{it} + \sigma \varepsilon_i + \gamma \varepsilon_{it}) e^{-\frac{\varepsilon_{it}^2}{2}} d\varepsilon_{it} \\
 &= \frac{1}{\pi} \int_{-\infty}^{\infty} e^{-u^2} du \int_{-\infty}^{\infty} \exp(\boldsymbol{\beta} \mathbf{x}'_{it} + \sqrt{2}\sigma u + \sqrt{2}\gamma v) e^{-v^2} dv \\
 &= \frac{1}{\pi} \sum_{h=1}^H \varphi_h \sum_{k=1}^K \omega_k \exp(\boldsymbol{\beta} \mathbf{x}'_{it} + \sqrt{2}\sigma u_h + \sqrt{2}\gamma v_k)
 \end{aligned}$$

## Appendix 2.5 Results Using Alternative Incentives

Table A2.1 Parameter Estimates for Different Types of Engagements (Pay Rate)

	Selection			Outcome		
	likes	comments	retweets	likes	comments	retweets
<b>Exclusion Variable</b>						
isEligible	2.543***	2.543***	2.547***			
<b>Incentive</b>						
payRate=High	0.018	0.020	0.018	-0.163	-0.138	-0.252
<b>Social Media Fanbase</b>						
log(followers)	0.131***	0.125***	0.123***	0.198**	-0.011	-0.060
verifiedRatio	0.550***	0.613***	0.587***	0.642	0.631	0.066
<b>Prior Activity Level</b>						
log(tweetNum)	0.044*	0.054**	0.058**	-0.005	-0.470***	-0.396***
log(taskNum)	0.535***	0.546***	0.543***	-0.389***	-0.464***	-0.535***
approvalRate	-0.030	-0.024	-0.026	0.828	1.522*	-0.737
<b>Community Embeddedness</b>						
regDays	-4.290***	-4.277***	-4.303***	0.092	2.048.	3.594***
log(friends)	-0.106**	-0.094*	-0.088*	0.502**	0.157	0.238
<b>Others</b>						
log(referralRwd)	-0.016.	-0.023*	-0.020*	0.011	0.062	-0.099
times	-0.151**	-0.146**	-0.157***	-0.280	-0.167	0.044
<b>Heterogeneity</b>						
$\delta$ (selection)				1.358***	1.349***	1.364***
$\sigma$ (outcome)				1.550***	2.232***	2.853***
$\gamma$ (outcome)				0.109	1.144***	1.423***
<b>Correlation</b>						
$\rho$ (endorser)				-0.228***	-0.206***	-0.306***
$\tau$ (endorser-task)				0.153	0.223	0.096
<b>Fitness</b>						
Log Likelihood				-7267.5	-7347.8	-7260.5
AIC				14650.9	14811.7	14637.0
BIC				15250.4	15411.1	15236.5

Table A2.2 Parameter Estimates for Different Types of Engagements (Linear Reward)

	Selection			Outcome		
	likes	comments	retweets	likes	comments	retweets
<b>Exclusion Variable</b>						
isEligible	2.535***	2.532***	2.528***			
<b>Incentive</b>						
actRwd	0.011*	0.012*	0.011.	0.015	0.063	0.114***
<b>Social Media Fanbase</b>						
log(followers)	0.123***	0.120***	0.121***	0.172*	-0.116	-0.150.
verifiedRatio	0.507***	0.552***	0.559***	0.593	0.628	-0.051
<b>Prior Activity Level</b>						
log(tweetNum)	0.045*	0.061**	0.048*	-0.012	-0.408***	-0.482***
log(taskNum)	0.535***	0.547***	0.549***	-0.356***	-0.358***	-0.674***
approvalRate	-0.023	-0.033	-0.016	0.764	1.403.	-0.022
<b>Community Embeddedness</b>						
regDays	-4.284***	-4.303***	-4.251***	-0.050	2.244*	3.048*
log(friends)	-0.105**	-0.107**	-0.100*	0.452**	0.099	0.259
<b>Others</b>						
log(referralRwd)	-0.016.	-0.022*	-0.023*	0.014	0.035	-0.043
times	-0.152**	-0.156***	-0.152***	-0.281	-0.169	0.008
<b>Heterogeneity</b>						
$\delta$ (selection)				1.358***	1.360***	1.359***
$\sigma$ (outcome)				1.564***	2.105***	2.795***
$\gamma$ (outcome)				0.005	1.040***	1.412**
<b>Correlation</b>						
$\rho$ (endorser)				-0.215***	-0.297***	-0.252***
$\tau$ (endorser-task)				-0.104	0.319*	0.269*
<b>Fitness</b>						
Log Likelihood				-7265.3	-7343.7	-7256.6
AIC				14646.6	14803.3	14629.2
BIC				15246.0	15402.8	15228.6

Table A2.3 Parameter Estimates for Different Types of Engagements (Linear GainLoss)

	Selection			Outcome		
	likes	comments	retweets	likes	comments	retweets
<b>Exclusion Variable</b>						
isEligible	2.544***	2.539***	2.539***			
<b>Incentive</b>						
gain	-0.018	-0.018	-0.018	-0.004	0.090	-0.064
loss	-0.253***	-0.245***	-0.237***	0.108	0.187	0.033
avgRwd	0.107*	0.102**	0.102*	0.044	0.048	0.250**
<b>Social Media Fanbase</b>						
log(followers)	0.113***	0.107***	0.107***	0.157*	-0.189*	-0.155.
verifiedRatio	0.411***	0.457***	0.450***	0.775	0.643	-0.191
<b>Prior Activity Level</b>						
log(tweetNum)	0.042*	0.051*	0.050*	0.026	-0.257*	-0.480***
log(taskNum)	0.539***	0.554***	0.538***	-0.366***	-0.453***	-0.386**
approvalRate	-0.004	-0.008	0.013	0.601	1.340.	-0.892
<b>Community Embeddedness</b>						
regDays	-4.344***	-4.375***	-4.258***	0.107	2.166.	3.134*
log(friends)	-0.098*	-0.099*	-0.087*	0.399*	-0.059	0.205
<b>Others</b>						
log(referralRwd)	-0.015.	-0.021*	-0.021*	0.026	0.027	0.056
times	-0.154***	-0.152***	-0.152**	-0.288	-0.218	0.015
<b>Heterogeneity</b>						
$\delta$ (selection)				1.353***	1.359***	1.339***
$\sigma$ (outcome)				1.578***	2.240***	2.846***
$\gamma$ (outcome)				0	1.175***	1.447***
<b>Correlation</b>						
$\rho$ (endorser)				-0.240***	-0.300***	-0.242***
$\tau$ (endorser-task)				0.005	0.238	0.207
<b>Fitness</b>						
Log Likelihood				-7251.3	-7329.4	-7242.0
AIC				14626.5	14782.8	14607.9
BIC				15267.3	15423.6	15248.7

The negative effects of follower number on comments and retweets in the outcome equation, though not well-supported in other models, are consistent with our reasoning that network size may have a negative effect on overall influence in high-effort behaviors (Katona et al. 2011).

Table A2.4 Parameter Estimates for Different Types of Engagements (Log Reward)

	Selection			Outcome		
	likes	comments	retweets	likes	comments	retweets
<b>Exclusion Variable</b>						
isEligible	0.967***	0.948***	0.952***			
<b>Financial Incentive</b>						
log(actRwd)	0.266***	0.267***	0.268***	-0.085	-0.010	-0.110
<b>Social Media Fanbase</b>						
log(followers)	0.017	0.007	0.008	0.202*	0.005	-0.081
verifiedRatio	0.154	0.198.	0.190.	0.646	0.806	0.386
<b>Prior Activity Level</b>						
log(tweetNum)	0.047*	0.056**	0.053**	0.006	-0.478***	-0.491***
log(taskNum)	0.543***	0.548***	0.542***	-0.354**	-0.414***	-0.311**
approvalRate	-0.003	0.005	0.015	0.727	1.638*	-0.462
<b>Community Embeddedness</b>						
regDays	-4.295***	-4.255***	-4.211***	-0.025	1.997.	3.130**
log(friends)	-0.100**	-0.093*	-0.080*	0.389*	0.046	-0.062
<b>Others</b>						
promIncm	-0.018*	-0.022*	-0.023*	0.018	0.030	-0.033
times	-0.142**	-0.144**	-0.147**	-0.298	-0.176	-0.127
<b>Heterogeneity</b>						
$\delta$ (selection)				1.360***	1.348***	1.344***
$\sigma$ (outcome)				1.597***	2.257***	2.661***
$\gamma$ (outcome)				0.031	1.124***	1.389***
<b>Correlation</b>						
$\rho$ (endorser)				-0.224***	-0.259***	-0.242***
$\tau$ (endorser-task)				-0.012	0.433**	0.380**
<b>Fitness</b>						
Log Likelihood				-7243.6	-7323.4	-7239.2
AIC				14603.2	14762.9	14594.4
BIC				15202.7	15362.3	15193.8

Due to the high correlation (0.98) between "isEligible" and "log(actRwd)", the parameter estimate on "isEligible" is very different from that in other models.

## Appendix 2.6 Additional Robustness Checks

Table A2.5 Effect of Eligibility on Effort Level

	# Words in retweets	Use of emoji
isEligible	-0.063	-0.137
log(gain)	0.010	0.039
log(loss)	0.007	0.071
log(avgRwd)	-0.130**	-0.105
log(followers)	0.083***	0.136.
verifiedRatio	0.139	0.493
log(tweetNum)	-0.020	0.083
log(taskNum)	0.106***	-0.114
approvalRate	0.120	0.455
regDays	-0.018	-1.207.
log(friends)	-0.030	0.168
log(referralRwd)	0.005	-0.012
times	0.015	-0.350
$\delta$		1.533***
$\sigma$	0.814***	
$\gamma$	0.448***	
Log Likelihood	-8406.1	-591.8

Since the number of words is a count variable and the use of emoji is binary, we use a Poisson Lognormal model with random effects on endorser level (similar to our outcome equation) and a Probit model with random effects on endorser level (similar to our selection equation) to estimate the effects of eligibility on these two effort metrics, respectively.

Table A2.6 Robustness to Model Complexity

	Full			Without $\tau$			Without $\tau$ & $\gamma$		
	likes	comments	retweets	likes	comments	retweets	likes	comments	retweets
<b>Selection</b>									
isEligible	2.724***	2.723***	2.720***	2.727***	2.730***	2.732***	2.732***	2.719***	2.723***
log(gain)	0.007	0.008	0.009	0.007	0.007	0.006	0.007	0.008	0.006
log(loss)	-0.073***	-0.074***	-0.073***	-0.074***	-0.076***	-0.075***	-0.074***	-0.072***	-0.073***
log(avgRwd)	0.034	0.036	0.037	0.039	0.038	0.036	0.040	0.024	0.029
log(followers)	0.083***	0.077***	0.079***	0.082***	0.077***	0.077***	0.082***	0.077***	0.077***
verifiedRatio	0.256*	0.297**	0.283**	0.251*	0.302**	0.299**	0.260*	0.290**	0.290**
log(tweetNum)	0.035.	0.045*	0.039.	0.035.	0.042*	0.036.	0.035.	0.041*	0.036.
log(taskNum)	0.555***	0.564***	0.561***	0.553***	0.560***	0.544***	0.554***	0.548***	0.556***
approvalRate	0.063	0.051	0.063	0.077	0.062	0.110	0.069	0.074	0.077
regDays	-4.375***	-4.41***	-4.358***	-4.374***	-4.399***	-4.324***	-4.400***	-4.345***	-4.337***
log(friends)	-0.088*	-0.084*	-0.086*	-0.086*	-0.079*	-0.064.	-0.087*	-0.072*	-0.069.
log(referralRwd)	-0.015.	-0.017.	-0.017.	-0.015.	-0.016.	-0.016.	-0.014	-0.017*	-0.018*
times	-0.142**	-0.142**	-0.147**	-0.142**	-0.139**	-0.139**	-0.144**	-0.144**	-0.146**
<b>Outcome</b>									
log(gain)	-0.039	-0.011	-0.061	-0.037	0.002	-0.088	-0.042	-0.072	-0.004
log(loss)	0.022	0.146	0.095	0.026	0.180.	0.067	0.015	0.162**	-0.001
log(avgRwd)	-0.092	0.174	0.333	-0.101	0.140	0.345	-0.065	-0.093	0.770***
log(followers)	0.221.	-0.039	-0.195	0.223.	-0.043	-0.177	0.216*	-0.028	-0.161*
verifiedRatio	0.652	1.105	0.989	0.650	1.097.	0.971	0.622	0.737.	0.000
log(tweetNum)	0.001	-0.478***	-0.252*	0.000	-0.396***	-0.260*	-0.014	-0.409***	-0.463***
log(taskNum)	-0.379***	-0.483***	-0.570***	-0.378***	-0.520***	-0.574***	-0.369***	-0.466***	-0.778***
approvalRate	0.704	1.708.	-0.300	0.706	1.707.	-0.338	0.672	-0.409	1.119*
regDays	0.102	2.826*	3.258**	0.101	2.784*	3.242**	0.124	3.692***	4.334***
log(friends)	0.420**	0.131	0.368	0.422**	0.065	0.358	0.419**	-0.432**	0.287.
log(referralRwd)	0.023	0.003	-0.095	0.021	-0.014	-0.092	0.021	0.101**	-0.204***
times	-0.285	-0.102	0.017	-0.285	-0.092	0.020	-0.276	0.102	1.009*
<b>Heterogeneity</b>									
$\delta$ (selection)	1.310***	1.315***	1.314***	1.31***	1.308***	1.277***	1.311***	1.290***	1.291***
$\sigma$ (outcome)	1.609***	2.190***	2.693***	1.607***	2.132***	2.616***	1.592***	2.106***	2.370***
$\gamma$ (outcome)	0.097	1.131***	1.371**	0.094	1.123***	1.315***			
<b>Correlation</b>									
$\rho$ (endorser)	-0.216***	-0.253***	-0.247***	-0.212***	-0.211***	-0.145**	-0.213***	-0.117**	-0.235***
$\tau$ (endorser-task)	0.008	0.199	0.312.						
<b>Fitness</b>									
Log Likelihood	-7229.7	-7304.5	-7216.5	-7229.7	-7304.8	-7219.0	-7229.5	-7319.4	-7251.7
AIC	14583.4	14733.0	14557.0	14581.4	14731.5	14560.0	14579.1	14758.8	14623.4
BIC	15224.2	15373.8	15197.8	15211.8	15362.0	15190.4	15199.2	15378.9	15243.5

Table A2.7 Robustness to Potential Outliers

	All Tasks			Without Task 2			Without Tasks 1-4		
	likes	comments	retweets	likes	comments	retweets	likes	comments	retweets
<b>Selection</b>									
isEligible	2.724***	2.723***	2.720***	2.776***	2.769***	2.769***	2.924***	2.920***	2.927***
log(gain)	0.007	0.008	0.009	0.001	0.003	0.002	0.005	0.006	0.004
log(loss)	-0.073***	-0.074***	-0.073***	-0.077***	-0.077***	-0.078***	-0.081***	-0.081***	-0.081***
log(avgRwd)	0.034	0.036	0.037	0.032	0.029	0.035	0.046	0.045	0.044
log(followers)	0.083***	0.077***	0.079***	0.085***	0.080***	0.079***	0.073***	0.071***	0.068***
verifiedRatio	0.256*	0.297**	0.283**	0.264*	0.313**	0.303**	0.174	0.185	0.183
log(tweetNum)	0.035.	0.045*	0.039.	0.034.	0.045*	0.041.	0.037.	0.040.	0.038.
log(taskNum)	0.555***	0.564***	0.561***	0.56***	0.572***	0.566***	0.596***	0.600***	0.589***
approvalRate	0.063	0.051	0.063	0.091	0.058	0.077	0.045	0.033	0.053
regDays	-4.375***	-4.410***	-4.358***	-4.433***	-4.479***	-4.444***	-4.632***	-4.594***	-4.535***
log(friends)	-0.088*	-0.084*	-0.086*	-0.085*	-0.083*	-0.075.	-0.089*	-0.084*	-0.078*
log(referralRwd)	-0.015.	-0.017.	-0.017.	-0.015.	-0.017.	-0.017.	-0.023*	-0.024**	-0.027**
times	-0.142**	-0.142**	-0.147**	-0.140**	-0.134**	-0.138**	-0.076	-0.077	-0.082.
<b>Outcome</b>									
log(gain)	-0.039	-0.011	-0.061	-0.035	-0.008	0.008	-0.110	-0.039	-0.021
log(loss)	0.022	0.146	0.095	0.030	0.200.	0.069	-0.046	0.094	0.150
log(avgRwd)	-0.092	0.174	0.333	-0.122	-0.053	0.344	0.029	0.041	0.319
log(followers)	0.221.	-0.039	-0.195	0.231.	0.013	-0.114	0.200	-0.123	-0.105
verifiedRatio	0.652	1.105	0.989	0.692	1.172.	0.557	0.499	0.718	0.749
log(tweetNum)	0.001	-0.478***	-0.252*	0.000	-0.46***	-0.455***	-0.013	-0.317***	-0.235
log(taskNum)	-0.379***	-0.483***	-0.570***	-0.378***	-0.436***	-0.432**	-0.343**	-0.414***	-0.513***
approvalRate	0.704	1.708.	-0.300	0.651	1.551	-0.400	0.941	1.169	-0.817
regDays	0.102	2.826*	3.258**	0.031	2.282.	3.317.	-0.140	1.095	2.066
log(friends)	0.42**	0.131	0.368	0.419**	-0.021	0.337	0.385*	-0.015	0.254
log(referralRwd)	0.023	0.003	-0.095	0.021	0.030	-0.113	0.030	0.048	-0.031
times	-0.285	-0.102	0.017	-0.287	-0.111	0.059	-0.243	-0.350	-0.176
<b>Heterogeneity</b>									
$\delta$ (selection)	1.31***	1.315***	1.314***	1.322***	1.332***	1.328***	1.368***	1.363***	1.346***
$\sigma$ (outcome)	1.609***	2.190***	2.693***	1.605***	2.282***	2.828***	1.600***	1.663***	2.210***
$\gamma$ (outcome)	0.097	1.131***	1.371**	0.124	1.186***	1.438**	0.000	1.620***	1.586***
<b>Correlation</b>									
$\rho$ (endorser)	-0.216***	-0.253***	-0.247***	-0.213***	-0.267***	-0.260***	-0.194**	-0.213**	-0.138*
$\tau$ (endorser-task)	0.008	0.199	0.312.	0.022	0.206	0.344*	0.000	0.058	0.094
<b>Fitness</b>									
Log Likelihood	-7229.7	-7304.5	-7216.5	-7168.4	-7234.5	-7147.1	-6543.4	-6583.6	-6497.7
AIC	14583.4	14733.0	14557.0	14460.7	14593.0	14418.3	13202.8	13283.2	13111.5
BIC	15224.2	15373.8	15197.8	15099.6	15231.9	15057.2	13794.9	13875.3	13703.5

## BIBLIOGRAPHY

- Abramowitz, M., and Stegun, I.A. 1972. *Handbook of Mathematical Functions*. Dover New York.
- Alexandrov, A., Lilly, B., and Babakus, E. 2013. "The Effects of Social- and Self-Motives on the Intentions to Share Positive and Negative Word of Mouth," *Journal of the Academy of Marketing Science* (41:5), pp. 531-546.
- Allison, P. 2002. "Bias in Fixed-Effects Cox Regression with Dummy Variables." Department of Sociology, University of Pennsylvania.
- Aral, S., Muchnik, L., and Sundararajan, A. 2009. "Distinguishing Influence-Based Contagion from Homophily-Driven Diffusion in Dynamic Networks," *Proceedings of the National Academy of Sciences* (106:51), pp. 21544-21549.
- Aral, S., and Van Alstyne, M. 2011. "The Diversity-Bandwidth Trade-Off," *American Journal of Sociology* (117:1), pp. 90-171.
- Aral, S., and Walker, D. 2011. "Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks," *Management Science* (57:9), pp. 1623-1639.
- Aral, S., and Walker, D. 2012. "Identifying Influential and Susceptible Members of Social Networks," *Science* (337:6092), pp. 337-341.
- Aral, S., and Walker, D. 2014. "Tie Strength, Embeddedness, and Social Influence: A Large-Scale Networked Experiment," *Management Science* (60:6), pp. 1352 - 1370.
- Bakshy, E., Hofman, J.M., Mason, W.A., and Watts, D.J. 2011. "Everyone's an Influencer: Quantifying Influence on Twitter," *Proceedings of the fourth ACM international conference on Web search and data mining: ACM*, pp. 65-74.
- Bakshy, E., Rosenn, I., Marlow, C., and Adamic, L. 2012. "The Role of Social Networks in Information Diffusion," *Proceedings of the 21st international conference on World Wide Web: ACM*, pp. 519-528.
- Bapna, R., Gupta, A., Rice, S., and Sundararajan, A. 2015. "Trust, Reciprocity and the Strength of Social Ties: An Online Social Network Based Field Experiment," *Working Paper*.
- Bapna, R., and Umyarov, A. 2015. "Do Your Online Friends Make You Pay? A Randomized Field Experiment on Peer Influence in Online Social Networks," *Management Science* (61:8), pp. 1902 - 1920.
- Barasch, A., and Berger, J. 2014. "Broadcasting and Narrowcasting: How Audience Size Affects What People Share," *Journal of Marketing Research* (51:3), pp. 286-299.
- Berndt, E.K., Hall, B.H., and Hall, R.E. 1974. "Estimation and Inference in Nonlinear Structural Models," *Annals of Economic and Social Measurement* (3:4), pp. 653-665.
- Bock, G.-W., Zmud, R.W., Kim, Y.-G., and Lee, J.-N. 2005. "Behavioral Intention Formation in Knowledge Sharing: Examining the Roles of Extrinsic Motivators, Social-Psychological Forces, and Organizational Climate," *Management Information Systems Quarterly* (29:1), pp. 87-111.
- Braun, M., and Moe, W.W. 2013. "Online Advertising Campaigns: Modeling the Effects of Multiple Ad Creatives," *Marketing Science* (32:5), pp. 753-767.
- Brewer, M.B. 1976. "Randomized Invitations: One Solution to the Problem of Voluntary Treatment Selection in Program Evaluation Research," *Social Science Research* (5:3), pp. 315-323.

- Burke, M. 2011. "Reading, Writing, Relationships: The Impact of Social Network Sites on Relationships and Well-Being," in: *Human-Computer Interaction Institute*. Carnegie Mellon University.
- Burt, D.R. 2001. "Bandwidth and Echo: Trust, Information, and Gossip in Social Networks," in *Networks and Markets: Contributions from Economics and Sociology*. Russell Sage Foundation.
- Butler, J.S., and Moffitt, R. 1982. "A Computationally Efficient Quadrature Procedure for the One-Factor Multinomial Probit Model," *Econometrica* (50:3), pp. 761-764.
- Byrd, R.H., Lu, P., Nocedal, J., and Zhu, C. 1995. "A Limited Memory Algorithm for Bound Constrained Optimization," *SIAM Journal on Scientific Computing* (16:5), pp. 1190-1208.
- Camarero, C., and San José, R. 2011. "Social and Attitudinal Determinants of Viral Marketing Dynamics," *Computers in Human Behavior* (27:6), pp. 2292-2300.
- Cantor, D., O'Hare, B.C., and O'Connor, K.S. 2008. "The Use of Monetary Incentives to Reduce Nonresponse in Random Digit Dial Telephone Surveys," in *Advances in Telephone Survey Methodology*. pp. 471-498.
- Centola, D. 2010. "The Spread of Behavior in an Online Social Network Experiment," *Science* (329:5996), pp. 1194-1197.
- Cheema, A., and Kaikati, A.M. 2010. "The Effect of Need for Uniqueness on Word of Mouth," *Journal of Marketing Research* (47:3), pp. 553-563.
- Chu, S.-C., and Kim, Y. 2011. "Determinants of Consumer Engagement in Electronic Word-of-Mouth (Ewom) in Social Networking Sites," *International Journal of Advertising* (30:1), pp. 47-75.
- Cohn, A., Fehr, E., and Goette, L. 2015. "Fair Wages and Effort Provision: Combining Evidence from a Choice Experiment and a Field Experiment," *Management Science* (61:8), pp. 1777-1794.
- Cox, D.R. 1972. "Regression Models and Life Tables," *Journal of the Royal Statistical Society. Series B* (34:2), pp. 187-220.
- De Bruyn, A., and Lilien, G.L. 2008. "A Multi-Stage Model of Word-of-Mouth Influence through Viral Marketing," *International Journal of Research in Marketing* (25:3), pp. 151-163.
- Duflo, E., and Saez, E. 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment," *Quarterly Journal of Economics* (118:3), pp. 815-842.
- Easley, D., and Kleinberg, J. 2010. *Networks, Crowds, and Markets*. Cambridge University Press.
- Goel, S., Anderson, A., Hofman, J., and Watts, D.J. 2015. "The Structural Virality of Online Diffusion," *Management Science* (62:1), pp. 180 - 196.
- Goel, S., Watts, D.J., and Goldstein, D.G. 2012. "The Structure of Online Diffusion Networks," in: *Proceedings of the 13th ACM Conference on Electronic Commerce*. Valencia, Spain: ACM, pp. 623-638.
- Granovetter, M. 1985. "Economic Action and Social Structure: The Problem of Embeddedness," *American Journal of Sociology* (91:3), pp. 481-510.
- Granovetter, M., and Soong, R. 1986. "Threshold Models of Interpersonal Effects in Consumer Demand," *Journal of Economic Behavior & Organization* (7:1), pp. 83-99.

- Granovetter, M.S. 1973. "The Strength of Weak Ties," *American Journal of Sociology* (78:6), pp. 1360-1380.
- Greene, W. 2009. "Models for Count Data with Endogenous Participation," *Empirical Economics* (36:1), pp. 133-173.
- Greenleaf, E.A. 1995. "The Impact of Reference Price Effects on the Profitability of Price Promotions," *Marketing Science* (14:1), pp. 82-104.
- Grier, S.A., and Deshpandé, R. 2001. "Social Dimensions of Consumer Distinctiveness: The Influence of Social Status on Group Identity and Advertising Persuasion," *Journal of Marketing Research* (38:2), pp. 216-224.
- Haenlein, M. 2013. "Social Interactions in Customer Churn Decisions: The Impact of Relationship Directionality," *International Journal of Research in Marketing* (30:3), pp. 236-248.
- Hall, D.B. 2000. "Zero-Inflated Poisson and Binomial Regression with Random Effects: A Case Study," *Biometrics* (56:4), pp. 1030-1039.
- Hall, J.A., and Valente, T.W. 2007. "Adolescent Smoking Networks: The Effects of Influence and Selection on Future Smoking," *Addictive behaviors* (32:12), pp. 3054-3059.
- Hardie, B.G., Johnson, E.J., and Fader, P.S. 1993. "Modeling Loss Aversion and Reference Dependence Effects on Brand Choice," *Marketing Science* (12:4), pp. 378-394.
- Heckman, J.J. 1979. "Sample Selection Bias as a Specification Error," *Econometrica* (47:1), pp. 153-161.
- Ho, J.Y., and Dempsey, M. 2010. "Viral Marketing: Motivations to Forward Online Content," *Journal of Business Research* (63:9), pp. 1000-1006.
- Hoff, P.D. 2005. "Bilinear Mixed-Effects Models for Dyadic Data," *Journal of the American Statistical Association* (100:469), pp. 286-295.
- Hu, Y., and Bulte, C.V.d. 2014. "Nonmonotonic Status Effects in New Product Adoption," *Marketing Science* (33:4), pp. 509-533.
- Ipeirotis, P.G., Provost, F., and Wang, J. 2010. "Quality Management on Amazon Mechanical Turk," in: *Proceedings of the ACM SIGKDD Workshop on Human Computation*. Washington DC: ACM, pp. 64-67.
- Iyengar, R., Bulte, C.V.d., and Lee, J.Y. 2015. "Social Contagion in New Product Trial and Repeat," *Marketing Science* (34:3), pp. 408-429.
- Iyengar, R., Van den Bulte, C., and Valente, T.W. 2011. "Opinion Leadership and Social Contagion in New Product Diffusion," *Marketing Science* (30:2), pp. 195-212.
- Jones, S.C. 1973. "Self-and Interpersonal Evaluations: Esteem Theories Versus Consistency Theories," *Psychological Bulletin* (79:3), p. 185.
- Kahneman, D., and Tversky, A. 1979. "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* (47:2), pp. 263-291.
- Kalyanaram, G., and Winer, R.S. 1995. "Empirical Generalizations from Reference Price Research," *Marketing Science* (14:3\_supplement), pp. G161-G169.
- Kang, J.-H., Lerman, K., and Getoor, L. 2013. "La-Lda: A Limited Attention Topic Model for Social Recommendation," in *Social Computing, Behavioral-Cultural Modeling and Prediction*. Springer, pp. 211-220.

- Katona, Z., Zubcsek, P.P., and Sarvary, M. 2011. "Network Effects and Personal Influences: The Diffusion of an Online Social Network," *Journal of Marketing Research* (48:3), pp. 425-443.
- Kempe, D., Kleinberg, J., and Tardos, É. 2003. "Maximizing the Spread of Influence through a Social Network," in: *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. Washington, D.C.: ACM, pp. 137-146.
- King, R.A., Racherla, P., and Bush, V.D. 2014. "What We Know and Don't Know About Online Word-of-Mouth: A Review and Synthesis of the Literature," *Journal of Interactive Marketing* (28:3), pp. 167-183.
- Kruskal, W. 1968. "When Are Gauss-Markov and Least Squares Estimators Identical? A Coordinate-Free Approach," *The Annals of Mathematical Statistics* (39:1), pp. 70-75.
- Lancaster, T. 2000. "The Incidental Parameter Problem since 1948," *Journal of Econometrics* (95:2), pp. 391-413.
- Langholz, B., and BORGAN, R. 1995. "Counter-Matching: A Stratified Nested Case-Control Sampling Method," *Biometrika* (82:1), pp. 69-79.
- Langholz, B., and Goldstein, L. 1996. "Risk Set Sampling in Epidemiologic Cohort Studies," *Statistical Science* (11:1), pp. 35-53.
- Lattin, J.M., and Bucklin, R.E. 1989. "Reference Effects of Price and Promotion on Brand Choice Behavior," *Journal of Marketing Research* (26:3), pp. 299-310.
- Leskovec, J., Adamic, L.A., and Huberman, B.A. 2007. "The Dynamics of Viral Marketing," *ACM Transactions on the Web* (1:1), p. 5.
- Long, X., and Nasiry, J. 2015. "Prospect Theory Explains Newsvendor Behavior: The Role of Reference Points," *Management Science* (61:12), pp. 3009 - 3012.
- Lovett, M.J., Peres, R., and Shachar, R. 2013. "On Brands and Word of Mouth," *Journal of Marketing Research* (50:4), pp. 427-444.
- Lu, Y., Jerath, K., and Singh, P.V. 2013. "The Emergence of Opinion Leaders in a Networked Online Community: A Dyadic Model with Time Dynamics and a Heuristic for Fast Estimation," *Management Science* (59:8), pp. 1783-1799.
- Ludford, P.J., Cosley, D., Frankowski, D., and Terveen, L. 2004. "Think Different: Increasing Online Community Participation Using Uniqueness and Group Dissimilarity," *Proceedings of the SIGCHI conference on Human factors in computing systems*: ACM, pp. 631-638.
- McPherson, M., Smith-Lovin, L., and Cook, J.M. 2001. "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology* (27), pp. 415-444.
- Minkler, M. 2012. *Community Organizing and Community Building for Health and Welfare*. Rutgers University Press.
- Moldoveanu, M.C., and Baum, J.A.C. 2011. "'I Think You Think I Think You're Lying': The Interactive Epistemology of Trust in Social Networks," *Management Science* (57:2), pp. 393-412.
- Narayan, V., and Yang, S. 2007. "Modeling the Formation of Dyadic Relationships between Consumers in Online Communities," *Available at SSRN: <http://ssrn.com/abstract=1027982>*.
- Nielsen. 2013. "Paid Social Media Advertising - Industry Update and Best Practices."
- Nitzan, I., and Libai, B. 2011. "Social Effects on Customer Retention," *Journal of Marketing* (75:6), pp. 24-38.

- Paolacci, G., Chandler, J., and Ipeirotis, P.G. 2010. "Running Experiments on Amazon Mechanical Turk," *Judgment and Decision making* (5:5), pp. 411-419.
- Porter, S.R., and Whitcomb, M.E. 2003. "The Impact of Contact Type on Web Survey Response Rates," *Public Opinion Quarterly* (67:4), pp. 579-588.
- Powers, D.E., and Swinton, S.S. 1984. "Effects of Self-Study for Coachable Test Item Types," *Journal of Educational Psychology* (76:2), pp. 266-278.
- Puhani, P. 2000. "The Heckman Correction for Sample Selection and Its Critique," *Journal of Economic Surveys* (14:1), pp. 53-68.
- Rand, W., and Rust, R.T. 2011. "Agent-Based Modeling in Marketing: Guidelines for Rigor," *International Journal of Research in Marketing* (28:3), pp. 181-193.
- Reagans, R., and McEvily, B. 2003. "Network Structure and Knowledge Transfer: The Effects of Cohesion and Range," *Administrative science quarterly* (48:2), pp. 240-267.
- Richardson, M., and Domingos, P. 2002. "Mining Knowledge-Sharing Sites for Viral Marketing," in: *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*. Edmonton, Alberta, Canada: ACM, pp. 61-70.
- Roberts, S.G.B., Dunbar, R.I.M., Pollet, T.V., and Kuppens, T. 2009. "Exploring Variation in Active Network Size: Constraints and Ego Characteristics," *Social Networks* (31:2), pp. 138-146.
- Schau, H.J., and Gilly, M.C. 2003. "We Are What We Post? Self-Presentation in Personal Web Space," *Journal of Consumer Research* (30:3), pp. 385-404.
- Schnabel, R.B., and Eskow, E. 1999. "A Revised Modified Cholesky Factorization Algorithm," *SIAM Journal on Optimization* (9:4), pp. 1135-1148.
- Schweidel, D.A., and Moe, W.W. 2014. "Listening in on Social Media: A Joint Model of Sentiment and Venue Format Choice," *Journal of Marketing Research* (51:4), pp. 387-402.
- Sharara, H., Rand, W., and Getoor, L. 2011. "Differential Adaptive Diffusion: Understanding Diversity and Learning Whom to Trust in Viral Marketing," *Proceedings of Fifth International AAAI Conference on Weblogs and Social Media*.
- Shi, Z., Rui, H., and Whinston, A.B. 2014. "Content Sharing in a Social Broadcasting Environment: Evidence from Twitter," *MIS quarterly* (38:1), pp. 123-142.
- Singer, E., and Ye, C. 2013. "The Use and Effects of Incentives in Surveys," *The Annals of the American Academy of Political and Social Science* (645:1), pp. 112-141.
- Snyder, C.R., and Fromkin, H.L. 1980. *Uniqueness: The Human Pursuit of Difference*. Plenum Press New York.
- Stephen, A.T., and Toubia, O. 2010. "Deriving Value from Social Commerce Networks," *Journal of Marketing Research* (47:2), pp. 215-228.
- Susarla, A., Oh, J.-H., and Tan, Y. 2012. "Social Networks and the Diffusion of User-Generated Content: Evidence from Youtube," *Information Systems Research* (23:1), pp. 23-41.
- Tajfel, H., and Turner, J.C. 1979. "An Integrative Theory of Intergroup Conflict," *The Social Psychology of Intergroup Relations* (33), p. 47.
- Therneau, T.M. 2000. *Modeling Survival Data: Extending the Cox Model*. Springer.
- Toubia, O., Goldenberg, J., and Garcia, R. 2014. "Improving Penetration Forecasts Using Social Interactions Data," *Management Science* (60:12), pp. 3049-3066.

- Toubia, O., and Stephen, A.T. 2013. "Intrinsic Vs. Image-Related Utility in Social Media: Why Do People Contribute Content to Twitter?," *Marketing Science* (32:3), pp. 368-392.
- Travis, H. 2002. *Causes of Delinquency*. Transaction Publishers.
- Trusov, M., Bodapati, A.V., and Bucklin, R.E. 2010. "Determining Influential Users in Internet Social Networks," *Journal of Marketing Research* (47:4), pp. 643-658.
- Trusov, M., Bucklin, R.E., and Pauwels, K. 2009. "Effects of Word-of-Mouth Versus Traditional Marketing: Findings from an Internet Social Networking Site," *Journal of Marketing* (73:5), pp. 90-102.
- Turner, J.C., and Oakes, P.J. 1986. "The Significance of the Social Identity Concept for Social Psychology with Reference to Individualism, Interactionism and Social Influence," *British Journal of Social Psychology* (25:3), pp. 237-252.
- Uzzi, B. 1997. "Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness," *Administrative Science Quarterly* (42:1), pp. 35-67.
- Watts, D.J., and Dodds, P.S. 2007. "Influentials, Networks, and Public Opinion Formation," *Journal of Consumer Research* (34:4), pp. 441-458.
- Weenig, M.W., and Midden, C.J. 1991. "Communication Network Influences on Information Diffusion and Persuasion," *Journal of Personality and Social Psychology* (61:5), p. 734.
- Weimann, G. 1983. "The Strength of Weak Conversational Ties in the Flow of Information and Influence," *Social Networks* (5:3), pp. 245-267.
- Wiatrowski, M.D., Griswold, D.B., and Roberts, M.K. 1981. "Social Control Theory and Delinquency," *American Sociological Review* (46:5), pp. 525-541.
- Winkelmann, R. 1998. "Count Data Models with Selectivity," *Econometric Reviews* (17:4), pp. 339-359.
- Wu, F., Huberman, B.A., Adamic, L.A., and Tyler, J.R. 2004. "Information Flow in Social Groups," *Physica A: Statistical Mechanics and its Applications* (337:1), pp. 327-335.
- Yoganarasimhan, H. 2012. "Impact of Social Network Structure on Content Propagation: A Study Using Youtube Data," *Quantitative Marketing and Economics* (10:1), pp. 111-150.
- Zeng, X., and Wei, L. 2013. "Social Ties and User Content Generation: Evidence from Flickr," *Information Systems Research* (24:1), pp. 71-87.
- Zhu, C., Byrd, R.H., Lu, P., and Nocedal, J. 1997. "Algorithm 778: L-BFGS-B: Fortran Subroutines for Large-Scale Bound-Constrained Optimization," *ACM Transactions on Mathematical Software* (23:4), pp. 550-560.