The Dissertation Committee for Qian Tang Certifies that this is the approved version of the following dissertation:

Essays on Social Media, Social Influence, and Social Comparison

Committee:

Andrew B. Whinston, Supervisor

Bin Gu

John Mote

Ashish Agarwal

Wen Wen

Essays on Social Media, Social Influence, and Social Comparison

by

Qian Tang, B.M.S.; M.M.; M.S. Econ.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

The University of Texas at Austin

August 2013

Dedicated to my parents.

Acknowledgements

I owe my deepest gratitude to my advisor, Dr. Andrew B. Whinston, who has taught me to conduct research projects that can make an impact and encouraged me to take challenges. I also would like to express my appreciation for the faculty members, Dr. Bin Gu, Dr. John Mote, Dr. Ashish Agarwal, and Dr. Wen Wen for their insightful comments that have helped provoke new thoughts on my research. It is my pleasure to thank the fellow doctoral students at the University of Texas at Austin for their support. Finally, this dissertation would not have been possible without the support of my family and friends. I would like to take this great opportunity to thank my husband, Brian Gu and my parents, Zhucai Tang and Guixiang Li, for their care, understanding, and encouragements along the way.

Essays on Social Media, Social Influence, and Social Comparison

Qian Tang, Ph.D.

The University of Texas at Austin, 2013

Supervisor: Andrew B. Whinston

Social networking and social media technologies have greatly changed the way information is created and transmitted. Social media has made content contribution an efficient approach for individual brand building. With abundant user generated content and social networks, content consumers are constantly subject to social influence. Such social influence can be further utilized to encourage pro-social behavior.

Chapter 1 examines the incentives for content contribution in social media. We propose that exposure and reputation are the major incentives for contributors. Besides, as more and more social media websites offer advertising-revenue sharing with some of their contributors, shared revenue provides an extra incentive for contributors who have joined revenue-sharing programs. We develop a dynamic structural model to identify a contributor's underlying utility function from observed contribution behavior. We recognize the dynamic nature of the content-contribution decision—that contributors are forward-looking, anticipating how their decisions impact future rewards. Using data collected from YouTube, we show that content contribution is driven by a contributor's desire for exposure, revenue sharing, and reputation and that the contributor makes decisions dynamically.

Chapter 2 examines how social influence impact individuals' content consumption decisions in social network. Specifically, we consider social learning and

network effects as two important mechanisms of social influence, in the context of YouTube. Rather than combining both social learning and network effects under the umbrella of social contagion or peer influence, we develop a theoretical model and empirically identify social learning and network effects separately. Using a unique data set from YouTube, we find that both mechanisms have statistically and economically significant effects on video views, and which mechanism dominates depends on the specific video type.

Chapter 3 studies incentive mechanism to improve users' pro-social behavior based on social comparison. In particular, we aim to motivate organizations to improve Internet security. We propose an approach to increase the incentives for addressing security problems through reputation concern and social comparison. Specifically, we process existing security vulnerability data, derive explicit relative security performance information, and disclose the information as feedback to organizations and the public. To test our approach, we conducted a field quasi-experiment for outgoing spam for 1,718 autonomous systems in eight countries. We found that the treatment group subject to information disclosure reduced outgoing spam approximately by 16%. Our results suggest that social information and social comparison can be effectively leveraged to encourage desirable behavior.

Table of Contents

List of Tables

List of Figures

Chapter 1: Content Contribution for Revenue Sharing and Reputation in Social Media

1.1 Introduction

Social media websites have become extremely popular recently for sharing contents, with 48-hour videos uploaded to YouTube every minute, 200 million tweets posted on Twitter every day, over 9,000 bloggers contributing on a wide range of topics on the Huffington Post every day, and over 6 billion photos shared on Flickr. Many very successful contributors have risen above the crowd. Michelle Phan is one such example. Starting as a no-name YouTuber in 2008, she has posted almost 200 videos and impressed viewers with her creative talent in makeup. She became a YouTube makeup sensation and now works for Lancôme as its first video makeup artist. Other people whose successes were made possible by their social media appearances include Justin Bieber, Jessica Rose, Kate Upton, and Graeme Anthony.

Although it is evident that social media can provide an alternative success path, why people contribute to social media and what they receive from content contribution is still under debate. In practice, after AOL purchased the Huffington Post for $315 million in 2011, thousands of unpaid bloggers filed a lawsuit against the Huffington Post, demanding compensation for their blog posts. However, the Huffington Post insisted that the bloggers' decisions to contribute for free were rational because they earned exposure and opportunities in return. Unlike the Huffington Post, certain websites such as YouTube, About.com, Break.com, and Epinions.com are paying some users for their contributions in the form of advertising-revenue sharing. In academia, few researchers have looked into the rationale behind content contribution in social media. Existing studies have found that knowledge contribution is driven by professional reputation, interpersonal ties, direct reciprocity, and enjoyment of helping others (Wasko and Faraj 2005). Investigations of open-source software development also showed that users voluntarily work on open-source tasks because of enjoyment of the work, direct need for software improvement, and the reputation that follows from making high-quality contributions (Fang and Neufeld 2009, Lakhani and Hippel 2003, Lerner and Tirole 2002). While these studies indicate potential drivers for content contribution, significant differences exist between social media and knowledge sharing or open-source software communities. In addition, extant research is based mainly on survey results that often ignore the dynamic nature of contribution.

In response to this gap, we develop an empirical framework to analyze the incentives for content contribution in social media. Relative to the extant literature, our research is distinguished in the following respects. First, we model the monetary payback for content contribution explicitly and examine how it motivates the contributors relative to other motivations, including exposure and reputation. Second, we allow the contributors to be forward-looking or myopic and test which attitude can better explain our observations. If contributors are forward-looking when making contributions, they would consider both current and future utilities. If they are myopic, they would consider only current utilities. Third, we model the changes in exposure and reputation dynamically so that current exposure and reputation are allowed to impact future exposure and reputation.

Using data from YouTube, the largest online community for user-generated videos, we find that instead of being myopic, contributors are forward looking, taking into consideration future benefits when making contribution decisions. Our results show that besides exposure, revenue sharing and reputation are two major incentives for content contribution. For contributors who have joined the revenue-sharing program, shared revenue has greatly increased their valuation of the views of their contents, which determines the advertising-revenue income. We also find it is better to measure a contributor's reputation in terms of relative performance than absolute achievement because of competition among contributors. More specifically, a contributor values relative rank among contributors with similar status more than relative rank among all contributors. These findings have important implications for both practitioners and researchers.

Our research contributes to the studies on social media in the following ways. First, this is one of the first empirical studies to examine and quantify various incentives for content contribution. While most empirical studies on social media focus on behavioral influences or interactions to determine statistical associations, very few look at the underlying incentives that fundamentally define such social behaviors (Manski 2000). Second, we use a dynamic structural model to identify the underlying parameters in contributors' utility function, which define their valuations for different incentives. Our approach has a number of advantages. A dynamic structural model allows us to capture individual behavior where individuals are forward looking and consider not only the current-period benefit of an action but also its future benefits. The model fits well with our research context as contribution decisions to social media impact not only a contributor's current-period utility but also future utilities. Our analysis confirms that contributors are indeed forward-looking when making contribution decisions. A dynamic structural model also allows us to recover underlying parameters of individual utility function from nonexperimental data based on the principle of revealed preference. The identification of the parameters enables counterfactual analysis by simulating the effect of changes in the underlying environment (e.g. an increase in revenue-sharing percentage), and to compare the explanatory power of competing theories (Reiss and Wolak 2005).

The remainder of this chapter is organized as follows. Section 1.2 is a review on related research. In Section 1.3, we provide an introduction of our research context. Section 1.4 describes the data. Section 1.5 presents the empirical approach with a proposed model for content contribution. Section 1.6 estimates the model with data and offers the estimation results. Section 1.7 discusses research findings, implications, and possible extensions.

1.2 Literature Review

Social media, also referred to as user-generated media, are new sources of online information created, initiated, circulated, and used by consumers intent on educating each other about products, brands, services, personalities, and issues (Dellarocas et al. 2010, Goldenberg et al. 2012). Users have both social presence and self-presentation on social-media platforms. On the one hand, according to social-presence theory (Short et al. 1976), the higher the social presence, the larger the social influence the communication partners have on each other's behavior. Self-presentation, on the other hand, can be used to influence others by controlling the impressions others form of them. In many cases, both social presence and self-presentation are used to influence others to gain rewards. Zhang and Zhu (2011) further showed the social benefits a contributor receives increase with group size. In addition, psychological reasons also help explain individuals' repeated participation in social media. Bateman et al. (2011) found that members might have psychological bonds to a community based on need, affect, and obligation. In particular, they found need-based commitment predicting thread reading, affect-based commitment predicting reply posting and moderating behaviors, and obligation-based commitment predicting only moderating behavior. Ma and Agarwal (2007) found that perceived identity verification can promote knowledge contribution.

Reputation has long been recognized as a key driver of content contribution in social media. Reputation could lead to social rewards such as approval, status, and respect (Jones et al. 1997). Given the enormous amount of information in social media, content consumers often face high levels of information asymmetry and information overload. In these situations, they tend to choose content from reputable contributors (Dellarocas 2003). Many social-media websites provide specially designed reputation mechanisms to facilitate online interactions (Dellarocas 2003, Rice 2012). A contributor's online reputation can also extend to that person's professional life (Stewart 2005), as more and more potential employers are taking a job applicant's online reputation into consideration. Thus, reputation is especially important for content contributors with career concerns or business motivations (Holmstrom 1999). However, there are no universal measures for reputation. Walker (2012) summarized that, although its measures are context specific, reputation is perception based, aggregated, and comparative, and that measures for reputation should reflect these three attributes.

For many social-media websites, reputation alone may not be strong enough to motivate content contribution, especially high-quality content contribution. The objective of social media websites such as YouTube, Facebook, and Twitter is to employ user-generated content to attract advertisers. Therefore, the market potential of contributed contents is important for these websites, whereas it is of lesser concern for the contributors. Consequently, many websites offer advertising-revenue sharing with contributors to resolve this misalignment in incentives. Besides social media, many other industries have also widely adopted revenue sharing to address similar principal-agent problems. Extant studies have analyzed the principal-agent problem in a variety of contexts and demonstrate revenue sharing could be an effective incentive mechanism to mitigate the problem (e.g., see Atkinson et al. (1988) in the context of a professional sports league; Black and Lynch (2004) in the context of workplace innovations; Arthur and Jelf (1999) in the context of union-management relations; and Mortimer (2008) in the context of business-to-business contract choices). The principal-agent problem is more prominent in online social media, as there exists no contractual relationship between the social media website and content contributors, and the quality of content is difficult to assess and monitor. Revenue sharing thus serves as an important strategic tool used by social media websites to motivate quality content contribution.

1.3 Research Context

We conduct our research based on YouTube, a video-sharing website founded in February 2005 by three former PayPal employees: Steve Chen, Chad Hurley, and Jawed Karim. The company is based in San Bruno, CA. In November 2006, Google bought YouTube for $1.65 billion, and since then it has operated as a subsidiary of Google. Ranked as the top online video website, YouTube allows billions of people to discover, watch, and share originally created videos. As of August 2011, it had drawn in 3.5 billion viewing sessions from 162 million unique viewers. Monetization of these views resulted in over $240 million of advertising revenue during 2010, according to Multichannel.com. Wall Street experts estimated that YouTube generated over $1.1 billion in annual revenue for Google in 2011.

Both unregistered and registered users can watch videos on YouTube, but only registered users can upload their videos. Registration is free and can be done in seconds. Each registered user has a unique YouTube page, called a YouTube channel. The channel presents a user's profile information, uploaded videos, and recent activities. Most videos on the website are generated by individual users. The uploaded videos can be up to 15 minutes in length, and that limit can be increased through account verification using a mobile phone. The videos' purposes can be entertaining, informational, or tutorial. The major video categories include auto and vehicles, comedy, education, entertainment, film and animation, gaming, how-to and style, news and politics, nonprofits and activism, people and blogs, pets and animals, science and technology, sports, and travel and events. Besides presenting information about the video contributor, a YouTube channel also allows viewers to subscribe to the contributor so they can get immediate notifications of new video postings. In this respect, subscribers can be considered loyal viewers of the contributor's videos. Although YouTube is a social broadcasting website, contributors can set their content to be private, or they can give only specific users permission to view. This kind of behavior is driven by a contributor's need to connect to or share with specific viewers. The specific incentive is beyond the scope of this study, and we do not consider such behavior as content contribution in this research.

YouTube makes money out of advertising with and within videos and channels. The ads can take the forms of standard video, expandable video, inVideo, inStream, rich media, and banner. YouTube enables inVideo ads only with partners and shares advertising revenue with them. To become a YouTube partner and obtain revenue sharing from video views, a contributor must regularly upload videos that are original and can generate thousands of views. A contributor can also publish extremely popular or commercially successful videos to apply for partial partnership on those specific videos only. YouTube does not publish specific requirements for partnership though, and reserves the right to make the final decisions.

YouTube is representative of other social-media websites and reflects a future trend in social media with respect to monetization of contents. Like YouTube, many social-media websites strive to monetize their content and user bases. Over 100 websites offer advertising-revenue sharing to their content contributors. For example, at About.com, Guides—who create contents to help users—are paid according to increase in page views in addition to a base payment, guaranteed to be $675 per month in the first two years and a minimum of $500 per month afterward. Shareapic.net pays its users $0.22 per 1,000 picture views. Epinions.com gives its users Eroyalties credits for writing reviews, and these credits are redeemable for U.S. dollars. Break.com pays its contributors up to $2,000 if their videos are selected to be posted on the homepage. The rewards from revenue sharing have made content contribution a full-time profession for many contributors. Therefore, what we study in the context of YouTube applies to these websites as well as other websites considering monetization. Our study helps the management of these social-media websites to gain a deep understanding of their contributors' valuations of exposure, monetary payback, and reputation. Such insight informs the design of incentive mechanisms such as whether revenue sharing is necessary to motivate contributions, how contributions change with shared revenue, and how many contributors to share revenue with.

1.4 Empirical Approach

This research aims to explore incentives for content contribution, including revenue sharing and reputation as well as intrinsic value from exposure. Therefore, our model needs to capture contributors' valuations of these factors. Although we do not observe their payoffs from revenue sharing or reputation directly, contributors' valuations can be inferred from changes in their contribution behavior according to the principle of revealed preference (Blundell et al. 2003).

1.4.1 Current-Period Utility Function

A contributor's current-period utility is influenced by four major factors: 1) exposure in the current period; 2) shared advertising revenue if the contributor is participating in revenue sharing; 3) accumulated reputation that may bring in psychological benefits such as the warm glow of doing something good, or potential economic benefits such as sponsorship, investment, or career opportunities in the current period; and 4) contribution costs that include the actual costs for producing[1] and posting videos and opportunity costs.

We model the current-period utility of exposure as a function of new views (v) and new subscribers (s) gained in the current period. New views and new subscribers refer to the number of video views and the number of new subscribers obtained in a given time period, respectively. We calculate new views and new subscribers from total views (V) and total subscribers (S) in two adjacent periods. Therefore, the intrinsic utility from increased exposure (U^1) for contributor i in period t can be written as

$$U_{it}^1 = \alpha_1 v_{it} + \alpha_2 s_{it} , \tag{1.1}$$

where $v_{it} = V_{it} - V_{it-1}$, and $s_{it} = S_{it} - S_{it-1}$.

Currently, on most social-media websites, not all contributors can share advertising revenue generated by their content. For example, About.com pays Guides only for their written content, and Break.com pays only contributors whose videos are selected to be posted on its homepage. YouTube is no exception, with revenue sharing limited to partners. We model the partnership status using a binary variable P_{it} where

$$\mathrm{P_{it}} = \begin{cases} 1, & \text{if provider } i \text{ is partner in period } t \\ 0, & \text{otherwise} \end{cases} . \tag{1.2}$$

[1] For many top contributors, production and posting decisions are separate and our analysis focuses on their posting decisions. For these contributors, production costs are sunken cost when they make posting decisions. Therefore, production costs are not part of contribution costs for these contributors.

Shared revenue depends on the views a partner's videos generate for the advertisements embedded in these videos. Although it is almost impossible (except for YouTube) to get the exact number of advertisement views, that figure is highly correlated with video views. Therefore, we measure the utility from revenue sharing (U^2) using an interaction term of partnership status and new views:

$$U_{it}^2 = \alpha_3 P_{it} v_{it} , \tag{1.3}$$

where α_3 measures the increased utility from new views because of revenue sharing.

A contributor's reputation may bring in many side opportunities such as investments, grants, sponsorship, and job offers. It also creates psychological benefits such as the warm-glow that makes a contributor feel good. Unlike revenue sharing, which has only a short-term impact on the contributor's utility, reputation is accumulated over time and has a long-lasting effect. Many studies suggest that reputation is a relative concept based on comparison with competitors or industry benchmark (Horner 2002). In social media, reputation has two aspects. One is popularity, which can be measured by total views (V). The other is quality, which can be approximated by total subscribers (S). One common criticism YouTube often receives is that it has too many funny or weird videos, which attract many viewers even though the videos are of very low quality. While it is not unusual for these videos to get many viewers, it is unlikely that these viewers would subscribe to the contributors' channels. One viral YouTube video, "Lily's Disneyland Surprise," received over 8 million views and made the contributor a fortune, but it only generated less than 2,000 subscribers for the contributor's channel. In this study, we use the rank of total views (RV) and the rank of total subscribers (RS) among all contributors to measure a contributor's reputation. Using ranks instead of absolute numbers also captures the competition on YouTube, or on social media in general, so that a contributor's reputation is not only determined by the contributor's achievement but also influenced by comparing the contributor's reputation with other contributors' performance. The utility from reputation (U^3) can thus be expressed as follows:

$$U_{it}^3 = R_{it} = \alpha_4 \frac{1}{RV_{it}} + \alpha_5 \frac{1}{RS_{it}}. \tag{1.4}$$

Since rank is an inverse measure of reputation, with a lower rank number meaning a higher reputation, we use the reciprocal of rank to keep a positive correlation between reputation and utility, and make coefficients α_4 and α_5 positive.

There is a certain cost associated with contribution, such as getting the video ready, adjusting the video to be consistent with YouTube requirements, and posting the video onto the YouTube channel. We assume that contribution cost (C) occurs only with actual contribution. For contributors who have their videos ready long before their posting decisions, the cost of producing video is sunk costs and the contribution costs thus include only costs related to the action of posting. So we have

$$\mathrm{C} = \alpha_6 a_{it}, \text{ where } a_{it} = \begin{cases} 1, & \text{making contribution} \\ 0, & \text{no contribution} \end{cases}, \tag{1.5}$$

with α_6 measuring the average contribution cost[2].

Meanwhile, a contributor's utility is affected by many other factors, such as the time availability, comments about previously uploaded videos, news and hot topics, and video-making capability. Unfortunately, such factors are unobservable to researchers. So we add an error term ε_{it} in utility function, which can be considered as the random utility shock in each period. The error term is allowed to be option specific (Arcidiacono and Miller 2011, Arellano 2000, Berry et al. 1995, Rust 1987), representing different unobservables for different options. We use $\varepsilon_{it}(a_{it})$ to model the component of utility

[2] The heterogeneity in contribution cost is absorbed in the random utility shocks. But since we cannot observe or measure the heterogeneity separately from the shocks, we cannot compare its magnitude with contribution benefits, or compare the heterogeneities of different provider groups.

of option a_{it} in time t which is known by the contributor but not by the researchers. $\varepsilon_{it}(0)$ captures unobserved utility for not contributing anything in period t. $\varepsilon_{it}(1)$ captures unobserved utility for making a contribution in period t. Both $\varepsilon_{it}(0)$ and $\varepsilon_{it}(1)$ are assumed to be Type I extreme values that are independent and identically distributed $(i.i.d.)$ across i and t. This assumption is commonly made to generate a Logit model for choice probabilities (Arcidiacono and Miller 2011, Arellano 2000, Rust 1987). The above discussion suggests that the current-period utility a contributor i receives in period t can be expressed as

$$U_{it} = \alpha_1 v_{it} + \alpha_2 s_{it} + \alpha_3 P_{it} v_{it} + \alpha_4 \frac{1}{RV_{it}} + \alpha_5 \frac{1}{RS_{it}} - \alpha_6 a_{it} + \varepsilon_{it}(a_{it}), \qquad (1.6)$$

1.4.2 Forecasting and Dynamic Transition

As we discussed above, the video provider makes a decision on a video-posting action (a_{it}) in each period. It is not realistic, however, to consider each of these decisions separately in each period since the influence of each decision would extend into the future. Specifically, different actions in the current period would lead to different numbers of new views, new subscribers, and reputations in the next period, and hence might result in different actions taken in the next period. Therefore, we characterize the provider's decisions in all periods as a sequential decision problem where a sequence of decisions must be made, with each decision affecting future decisions. In this dynamic setting, we need to model how contributors predict future states given their contribution decisions. The constructive procedure used for solving such problems is dynamic programming (Rust 1994). One important concept of dynamic programming is state variables, whose values completely specify the instantaneous situation of the process. The decisions influence total utilities by affecting change in state variables. Therefore, the state variables should include the variables influenced by the contributor's decisions, and thus influence the contributor's utility. According to (1.6), the state variables of our model should include at least $\mathrm{v_{it}}, \mathrm{s_{it}}$, P_{it}, $\mathrm{RV_{it}}$, and $\mathrm{RS_{it}}$, which determine current-period utility.

Contributors facing decision problems would anticipate the outcomes of choosing different actions (Manski 2000). The outcomes include changes in state variables and resulting utilities in future periods. According to Manski (2000), contributors forming expectations may seek to draw lessons from observations of the actions chosen and outcomes experienced both by themselves and others. Consequently, we model the contributor's expectation of the change in state variables using the pattern discovered empirically from observed data.

The changes of $\mathrm{v_{it}}, \mathrm{s_{it}}$, P_{it}, $\mathrm{RV_{it}}$, and $\mathrm{RS_{it}}$ over time are influenced by the provider's actions as well as other factors. The transition of s_{it} can be considered as a product-diffusion process, where the product is the contributor's YouTube channel. The most widely used new-product growth model is the Bass diffusion model (1969), which considers both mass media and interpersonal communication channels. The central proposition of the Bass diffusion model implies that

$$n(t) = pm + (q - p)N(t) - (q/m)[N(t)]^2, \tag{1.7}$$

where m represents potential market size, $n(t)$ is the number of new adopters at time t, $N(t)$ is cumulative number of adopters at time t, and p and q are coefficients for innovation and imitation. This model has been widely used in forecasting consumer demand. Based on this form of the Bass diffusion model, we model the transition of s_{it+1} as a function of S_{it} and S_{it}^2. In the specific context of YouTube, however, the increase in a contributor's channel subscribers also depends on the increase in a contributor's video views because new subscribers have to be generated out of new viewers. So we include v_{it+1} to capture the conversion from viewers into subscribers. We also expect to observe a reputation effect by which contributors with higher reputations can attract more new subscribers. Thus the reputation measures RV_{it} and RS_{it} may impact s_{it+1}. Some factors that also influence the evolution of new subscribers are specific to the provider, resulting in different evolution paths for different providers. These factors are unobservable to researchers. For example, better-looking providers in the videos may attract more viewers to subscribe to their channels. We denote such provider-specific variables as τ_i.

$$E(s_{it+1}) = \gamma_{20} + \gamma_{21}S_{it} + \gamma_{22}S_{it}^2 + \gamma_{23}v_{it+1} + \gamma_{24}v_{it+1}^2 + \gamma_{25}\frac{1}{RV_{it}} + \gamma_{26}\frac{1}{RS_{it}} + \tau_i. \quad (1.8)$$

Although providers can anticipate the expected value of s_{it+1} based on the diffusion path we described above (in (1.8), where E(.) denotes expectation), they do not have the perfect knowledge of the realization of s_{it+1}. The uncertainty lies in many aspects, such as competition from other channels or websites, ease of subscription, and exogenous change in total viewer base. It is thus impossible for providers to know the actual s_{it+1} beforehand. We assume that the realized s_{it+1} follows a normal distribution with mean $E(s_{it+1})$ as in (1.8) and variance σ_s^2. The uncertainty is summarized in a random subscription shock ς_{it+1}. So we have

$$s_{it+1} = \gamma_{20} + \gamma_{21}S_{it} + \gamma_{22}S_{it}^2 + \gamma_{23}v_{it+1} + \gamma_{24}v_{it+1}^2 + \gamma_{25}\frac{1}{RV_{it}} + \gamma_{26}\frac{1}{RS_{it}} + \tau_i + \varsigma_{it+1},$$

$$\varsigma_{it+1} \sim N(0, \sigma_\varsigma^2). \quad (1.9)$$

The expected new views v_{it+1} is essentially the aggregation of the diffusion processes of all contributors' videos. Existing studies have shown that the Bass diffusion model also applies to the demand for a product category and for individual brands (Ma and Agarwal 2007). Moreover, Hendricks and Sorensen (2000) demonstrated that a strong spillover effect exists among sales of music albums by the same artist. They find that the release of an artist's new album increases the sales of that same artist's old albums, especially if the new release is a hit, and vice versa. As a result, we assume that the Bass diffusion model also applies to the transition of v_{it+1}. What makes the process different from the traditional diffusion process is that YouTube's subscription function makes it very convenient for subscribers to check out new videos. Therefore, S_{it} and $a_{it}S_{it}$ are included to capture two different impacts of total subscribers on increased views. While S_{it} captures new views for the contributor's existing contents, $a_{it}S_{it}$ captures new views from existing subscribers if the contributor posts new videos. Essentially, λ_3 measures the herding effect caused by the existing subscriber base, and λ_4 measures the degree of loyalty of existing subscribers.

$$E(v_{it+1}) = \gamma_{10} + \gamma_{11}V_{it} + \gamma_{12}V_{it}^2 + \gamma_{13}S_{it} + \gamma_{14}a_{it}S_{it} + \gamma_{15}\frac{1}{RV_{it}} + \gamma_{16}\frac{1}{RS_{it}} + \mu_i \quad (1.10)$$

Similar to the evolution of expected s_{it+1}, we also include reputation effect from RV_{it} and RS_{it}, and a contributor-specific effect (μ_i). Because of the uncertainty of realized v_{it+1}, the evolution of $v_{it+1}$is also a stochastic process in which v_{it+1} follows a normal distribution with mean $E(v_{it+1})$ as in (1.10) and variance σ_v^2. The random view shock is denoted as ξ_{it+1}. So we have

$$v_{it+1} = \gamma_{10} + \gamma_{11}V_{it} + \gamma_{12}V_{it}^2 + \gamma_{13}S_{it} + \gamma_{14}a_{it}S_{it} + \gamma_{15}\frac{1}{RV_{it}} + \gamma_{16}\frac{1}{RS_{it}} + \mu_i + \xi_{it+1},$$

$$\xi_{it+1} \sim N\left(0, \sigma_\xi^2\right). \qquad (1.11)$$

A contributor's reputation presumably is highly correlated over time and endogenous so that it changes according to the contributor's performance in each period. In this study, we use relative ranks in video views and subscribers as measures for reputation to emphasize competition among contributors. Therefore, the uncertainties come from the unpredictable performance of all other contributors. According to knowledge-capital models (Doraszelski and Jaumandreu 2008, Griliches 2000), we consider both reputation in video views and reputation in subscribers to be governed by a first-order Markov process that

$$RV_{it+1} = \gamma_{30} + \gamma_{31} RV_{it} + \gamma_{32} v_{it+1} + \omega_{it+1}, \;\; w_{it+1} \sim N(0, \sigma_w^2) \qquad (1.12)$$

$$RS_{it+1} = \gamma_{40} + \gamma_{41} RS_{it} + \gamma_{42} s_{it+1} + r_{it+1}, \quad\; r_{it+1} \sim N(0, \sigma_r^2).$$

Theoretically, as video views and subscribers increase, the probability of a provider becoming a YouTube partner increases as well. YouTube can also terminate a provider's partnership if the provider is not in compliance with YouTube's terms of service. However, in reality, we don't observe any instance of a partner becoming a nonpartner. So we assume that once becoming a partner, the provider would remain a partner for all future periods. Therefore, we assume the probability of being a partner in the next period is conditional on current partnership status such that

$$\mathbf{Pr}(P_{it+1} = 1) = \begin{cases} 1, & if \; P_{it} = 1 \\ 1/(1 + \mathrm{e}^{-(\gamma_{50} + \gamma_{51} V_{it+1} + \gamma_{52} S_{it+1})}), & if \;\; P_{it} = 0 \end{cases}. \qquad (1.13)$$

To focus on an optimal Markovian decision rule where each contributor's action depends only on the current state (Rust 1987, 1994), we need to include all variables in the current period that influences future utilities either directly or indirectly in the state variables. So state variables (State_{it}) are constructed to include S_{it}, V_{it}, μ_i, and τ_i, as well as s_{it}, v_{it}, P_{it}, RV_{it}, RS_{it}, and ε_{it}, as in (1.14). S_{it} and V_{it} are updated determinately simply by adding the increase to the existing cumulative number. μ_i and τ_i are constant over time.

$$\mathrm{State}_{it} = (\, v_{it}, s_{it}, RS_{it}, RV_{it}, P_{it}, S_{it}, V_{it}, \mu_i, \tau_i, \varepsilon_{it}) = (x_{it}, \varepsilon_{it}) \,. \qquad (1.14)$$

The dynamic transition process can be described as

$$v_{it+1} = \gamma_{10} + \gamma_{11}V_{it} + \gamma_{12}V_{it}^2 + \gamma_{13}S_{it} + \gamma_{14}a_{it}S_{it} + \gamma_{15}\frac{1}{RV_{it}} + \gamma_{16}\frac{1}{RS_{it}} + \mu_i + \xi_{it+1}$$

$$s_{it+1} = \gamma_{20} + \gamma_{21}S_{it} + \gamma_{22}S_{it}^2 + \gamma_{23}v_{it+1} + \gamma_{24}v_{it+1}^2 + \gamma_{25}\frac{1}{RV_{it}} + \gamma_{26}\frac{1}{RS_{it}} + \tau_i + \varsigma_{it+1}$$

$$RV_{it+1} = \gamma_{30} + \gamma_{31}RV_{it} + \gamma_{32}v_{it+1} + \omega_{it+1}$$

$$RS_{it+1} = \gamma_{40} + \gamma_{41}RS_{it} + \gamma_{42}s_{it+1} + r_{it+1} \tag{1.15}$$

$$\mathbf{Pr}(P_{it+1} = 1) = \begin{cases} 1, & \text{if } P_{it} = 1 \\ 1/(1 + \mathrm{e}^{-(\gamma_{50}+\gamma_{51}V_{it+1}+\gamma_{52}S_{it+1})}), & \text{if } P_{it} = 0 \end{cases}$$

$$V_{it+1} = v_{it+1} + V_{it}$$

$$S_{it+1} = s_{it+1} + S_{it}$$

$$\mu_{it+1} = \mu_{it} = \mu_i$$

$$\tau_{it+1} = \tau_{it} = \tau_i$$

where $\xi_{it+1} \sim N(0, \sigma_\xi^2), \varsigma_{it+1} \sim N(0, \sigma_\varsigma^2), w_{it+1} \sim N(0, \sigma_w^2)$, and $r_{it+1} \sim N(0, \sigma_r^2)$.

Although μ_i, and τ_i are unobservable to researchers, they can be recovered from the data using panel-data methods (Bajari 2007). Because the evolution paths of v_{it+1}, s_{it+1}, RV_{it+1}, RS_{it+1}, and P_{it+1} are stochastic processes, we denote the Markov transition density for a state variable $(x_{it}, \varepsilon_{it})$ when action a_{it} is chosen as $\Pr(x_{it+1}, \varepsilon_{it+1} | a_{it}, x_{it}, \varepsilon_{it})$, representing the provider's subjective beliefs about an uncertain future state according to (1.15). x_{it} consists of variables observable either directly or indirectly by recovering them from other observables.

1.4.3 Markov Decision Process

Because of the uncertainties in realized next-period state variables, providers do not have perfect knowledge about future state variables when they make contribution decisions in each period. Although they are able to anticipate the expected values of future state variables, the actual values depend on the realized $\xi_{\mathrm{it}+1}$, $\varsigma_{\mathrm{it}+1}$, ω_{it+1} and r_{it+1} that occur after the contribution decision has been completely carried out. As a result, the contributor's decisions can change the state variables only probabilistically. The contributor's goal is to maximize the expected total utilities from all future periods, given the current state variables. We model these contributors as forward-looking instead of myopic because most of the contribution benefits are realized several periods after the contribution. Dynamic structural models have been used to capture the forward-looking feature of content contribution in social media. A dynamic structural model assumes that when making the decisions, contributors not only consider current-period utility but also take into account the discounted expected future utilities over an infinite time horizon. Therefore, the objective function can be described as

$$\max_{a_{it}} \mathbf{E}[\sum_{\tau=t}^{\infty} \beta^{\tau-t}\, U_{i\tau}\big(a_{i\tau}, x_{\mathrm{i}\tau}, \varepsilon_{i\tau}(a_{i\tau})\big|a_{it}, x_{\mathrm{it}}, \varepsilon_{it}(a_{it})\big)], \tag{1.16}$$

where β is the common discount factor. The myopia case is also included in the model simply by setting $\beta = 0$. (1.16) dynamically models a contributor's maximization problem, given the contributor's expectation of future utilities based on current state variables.

In each period, the contributor observes current-period state variables, makes a decision based on this information, and realizes current-period utility. In the next period, the state variables update and are revealed to the contributor, who then makes another decision based on the new state. We assume the contributor makes a decision at the end of each period. The model is equivalent, however, if the decision is made at the beginning of each period. The detailed decision process can be described as follows.

In period t

- First, a contributor observes the value of that contributor's current-state variables $x_{it} = (s_{it}, v_{it}, RV_{it}, RS_{it}, S_{it}, V_{it}, P_{it}, \mu_i, \tau_i, \varepsilon_{it})$ and believes these variables evolve according to 1.15.
- Second, because of uncertainties in ξ_{it+1}, ς_{it+1}, ω_{it+1}, r_{it+1} and P_{it+1}, the contributor forms expectations of $v_{it+1}, s_{it+1}, RV_{it+1}, RS_{it+1}$ and P_{it+1} in period t+1. The provider has two available options: $a_{it} = 0$ or 1. Since a_{it} influences v_{it+1}, v_{it+1} influences s_{it+1}, and v_{it+1} and s_{it+1} influence RV_{it+1} and RS_{it+1}, two options lead to different values of current-period utility U_{it} and different expected values of x_{it+1} in the next period. For each option, the contributor considers what to choose in the next period if the contributor is in the resulting new state $E(x_{it+1})$ with one fewer period remaining, and calculates expected utility from the next period, and so on.
- Eventually, the contributor adds up the discounted expected utilities from all future periods and current-period utility for each option, chooses the one that generates higher total utilities, and acts upon the chosen a_{it}. Meanwhile, the contribution cost is realized if $a_{it} = 1$.

In period t+1

- The contributor observes realized ξ_{it+1}, ς_{it+1}, ω_{it+1}, r_{it+1}, and P_{it+1}. The state evolves to realized $(s_{it+1}, v_{it+1}, RV_{it+1}, RS_{it+1}, S_{it+1}, V_{it+1}, P_{it+1}, \mu_i, \tau_i, \varepsilon_{it+1})$, which may be different from the provider's anticipation. Again, the contributor makes another decision of a_{it+1} based on information of x_{it+1} and imperfect knowledge about state evolution (1.15).

Given the setup of our model, we have an infinite-horizon, discounted Markovian decision problem, the solution of which is given by a stationary decision rule $a_{it} =$

$f(\mathrm{x_{it}}, \varepsilon_{\mathrm{it}})$ such that the contributor's optimal decision is a function of state variables. The optimal value function V_θ is a unique solution to the Bellman equation given by

$$V_\theta(x_{\mathrm{it}}, \varepsilon_{it}) = \mathbf{max}_{a_{it}\in\{0,1\}}\, u_{it}(a_{it}, x_{\mathrm{it}}) + \varepsilon_{it}(a_{it}) + \beta\mathbf{E}[V_\theta(x_{\mathrm{it+1}}, \varepsilon_{it+1})|a_{it}, x_{\mathrm{it}}, \varepsilon_{it}(a_{it})], \tag{1.17}$$

where $\mathbf{E}[V_\theta(x_{\mathrm{it+1}}, \varepsilon_{it+1})|a_{it}, x_{\mathrm{it}}, \varepsilon_{it}(a_{it})] =$

$$\int_{x_{\mathrm{it+1}}}\int_{\varepsilon_{it+1}} V_\theta(x_{\mathrm{it+1}}, \varepsilon_{it+1})\mathbf{P\,r}\big(dx_{\mathrm{it+1}}, d\varepsilon_{it+1}\big|a_{it}, x_{\mathrm{it}}, \varepsilon_{it}(a_{it})\big)$$

and the optimal decision rule is defined by

$$f(x_{\mathrm{it}}, \varepsilon_{it}) \equiv \underset{a_{it}\in\{0,1\}}{\mathbf{argmax}}\big[u_{it}(a_{it}, x_{\mathrm{it}}) + \varepsilon_{it}(a_{it}) + \beta\mathbf{E}[V_\theta(x_{\mathrm{it+1}}, \varepsilon_{it+1})|a_{it}, x_{\mathrm{it}}, \varepsilon_{it}(a_{it})]\big] \tag{1.18}$$

To simplify the numerical integration required in the Bellman equation (1.17), a conditional independence assumption is adopted from existing dynamic programming literature (Aguirreagabiria and Mira 2002, Holmstrom 1999, Rice 2012) such that

$$\mathbf{Pr}\big(x_{\mathrm{it+1}}, \varepsilon_{it+1}\big|a_{it}, x_{\mathrm{it}}, \varepsilon_{it}(a_{it})\big) = \mathbf{Pr}(x_{\mathrm{it+1}}|a_{it}, x_{\mathrm{it}})\,\mathbf{Pr}(\varepsilon_{it+1}|x_{\mathrm{it+1}}). \tag{1.19}$$

This conditional independence assumption allows us to simulate state evolution and random shock generation separately. Taking the standard approach for dynamic structural modeling, we set β to be constant (Aguirreagabiria and Mira 2007, Bajari et al. 2007), equal to 0.97. The reason for not being able to identify β is explained in detail in Section 6.2. However, we test the forward-looking model where $\beta = 0.97$, against the myopic model where $\beta = 0$. θ is the set of all the parameters to be estimated, and $\alpha = (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6)$ is the vector of structural parameters in the utility function.

1.5 Data

1.51 Data Collection

We collect a panel of data on two different groups of YouTube contributors for two months from January 9th, 2012, to March 10th, 2012. Set 1 includes the top 1,000 contributors of the most viewed YouTube channels in June 2011, while Set 2 consists of 2,236 new contributors who joined YouTube during January 2012.

The key decision a contributor makes in online social media is timing—when and how frequently to contribute content. To operationalize the decision process, we consider each contributor's decision in a given period. For parsimony, we treat multiple video postings in each period as one posting. Therefore, action in each period is whether a contributor posts any new videos. We define time period as one day and collect data daily. We choose a short time period because most YouTube videos have an extremely brief lifetime. Figure 1.1 plots the growth in views several sample videos received after being posted on YouTube and the average curves. The figure shows most videos received the greatest number of views in the first four to five days. After that time, views increased only marginally. We also collect information on two key statistics published by YouTube on the popularity of each contributor: 1) total video views (V_{it}) , which counts the number of times a contributor's videos have been viewed, and 2) total subscribers (S_{it}), which counts the number of viewers who have subscribed to the contributor's channel. For reputation measure, we collect data from vidstatsx.com on 1) video view rank (RV_{it}), the provider's rank among all contributors in terms of cumulative video views, and 2) subscriber rank (RS_{it}), the provider's rank among all contributors in terms of total subscribers. Because vidstatsx.com provides ranks only for the top 9,999 contributors, we record any missing rank as 10,000. For partnership status, we check YouTube pages for each contributor's videos. If any video has any inVideo ads, we treat the contributor as a partner ($P_{it} = 1$), and otherwise as a nonpartner ($P_{it} = 0$).

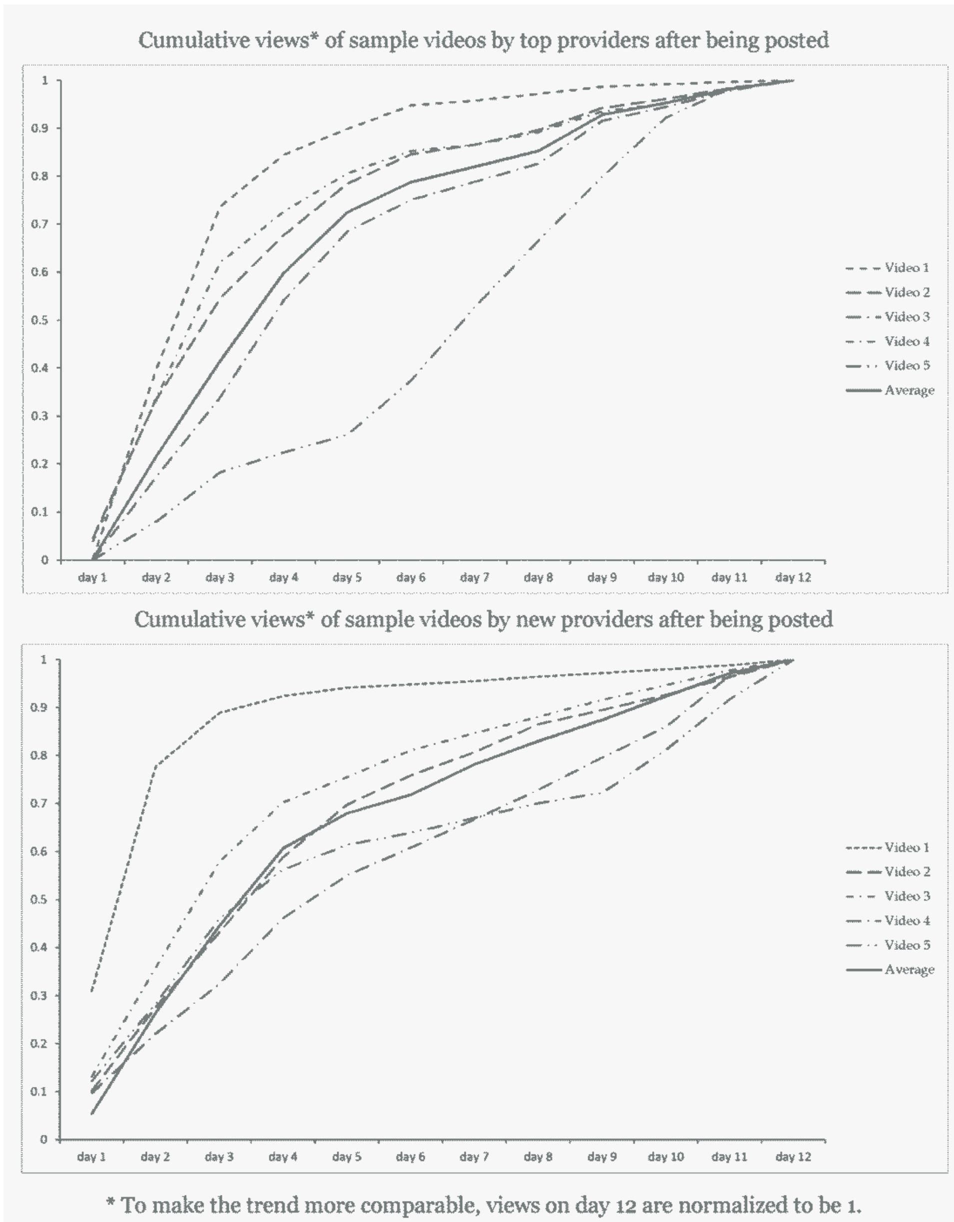

Figure 1.1: Cumulative Views of Sample Videos by Top Providers (top) and New Providers (bottom) after Being Posted

We collect data daily from YouTube and vidstatsx.com. For both groups, we delete the samples with missing values or no video postings during the data-collection period. The data-cleaning process reduced the sample size to 823 for Set 1 of the top contributors, and to 1252 for Set 2 of new contributors. For simplification, we refer to Set 1 as top sample set and to Set 2 as new sample set. While all the top contributors were partners during the data-collection period, all the new providers were nonpartners at the beginning, with 438 becoming partners and 814 remaining nonpartners during the period. We allow the size for a new sample set to be slightly larger than for a top sample set so we can get balanced partner versus nonpartner observations, since some new contributors became partners in between. In the estimation, we pool all the samples together as well as estimate partners and nonpartners separately. Figure 1.2 plots the distribution of the number of videos posted by sample providers during the data-collection period. It shows that top and new contributors posted a similar number of videos.

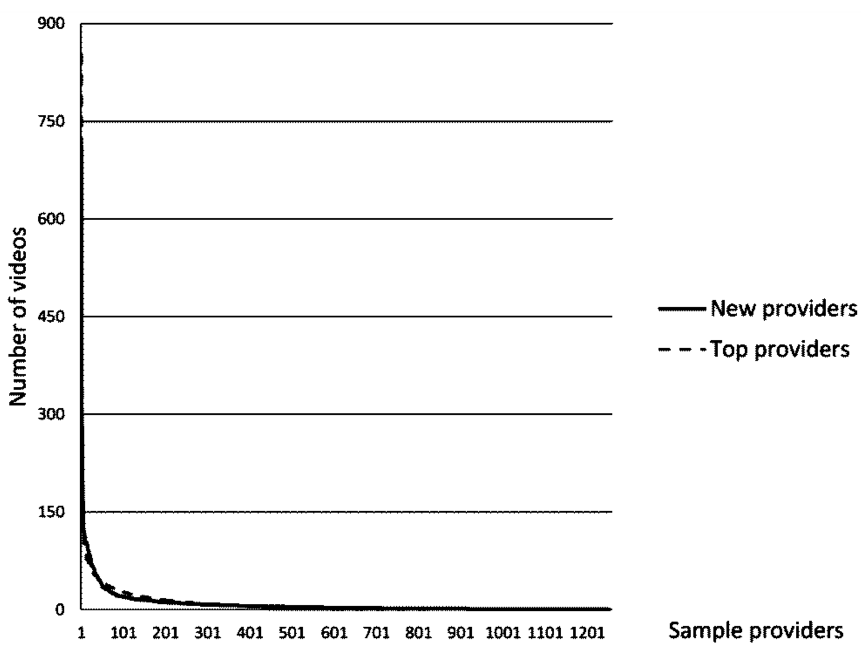

Figure 1.2: Distribution of Number of New Videos Posted

1.5.2 Summary Statistics

Table 1.1 provides summary statistics for the two sample sets. The statistics show that the two sets vary greatly from each other in terms of number of views, number of subscribers, views rank, subscribers rank, video-posting probability, partnership status, and number of videos. The top sample set accumulates far more video views, channel subscribers, and videos, and thus has a higher ranking in terms of video views and subscribers. The new sample set posts videos more frequently though, and some new contributors post even more videos (2,782) than the most productive contributor (with 989 videos posted) in the top sample set. While all the top providers are YouTube partners, most contributors in the new sample set are nonpartners.

Table 1.1: Descriptive Statistics

	Top sample set			
Variable	Mean	SD	Min	Max
video views	159,696,110	360,497,854	3,919,009	6,911,122,609
subscribers	364,196	493,522	12,294	5,297,145
views rank	4,631	4,196	1	10,000
subscribers rank	1,566	2,305	1	10,000
action (a)	0.1302	0.3366	0	1
partnership (P)	1	0	1	1
videos	213	217	1	989
sample size		823		
	New sample set			
Variable	Mean	SD	Min	Max
video views	180,821	863,195	0	23,045,802
subscribers	417	2,629	46	75,592
views rank	9,984	223	3,761	10,000
subscribers rank	9,980	341	1,949	10,000
action (a)	0.2019	0.4015	0	1
partnership (P)	0.1988	0.3991	0	1
videos	38	105	0	2,782
sample size		1,252		

The distribution of views and subscribers across contributors for the two sets are presented in Figures 1.3–1.6 (left). Data for both new sample sets are highly skewed. Therefore, we use the log transformation of views and subscribers to control for skewness in the data (Susarla et al. 2012). Figures 1.3–1.6 (right) plot the log-transformed views and subscribers. In the rest of the chapter, S and V refer to log-transformed total subscribers and video views, and s and v refer to the change in the log-transformed new subscribers and new views. Table 1.2 shows the summary of the key variables used in our estimation for the two sets pooled together.

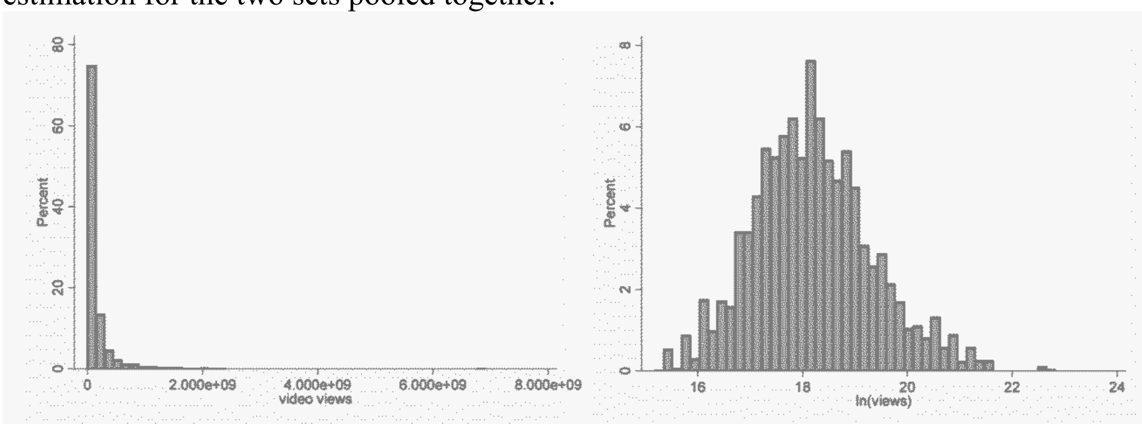

Figure 1.3: Distribution of Views (left) and ln(views) (right) for Top Sample Set

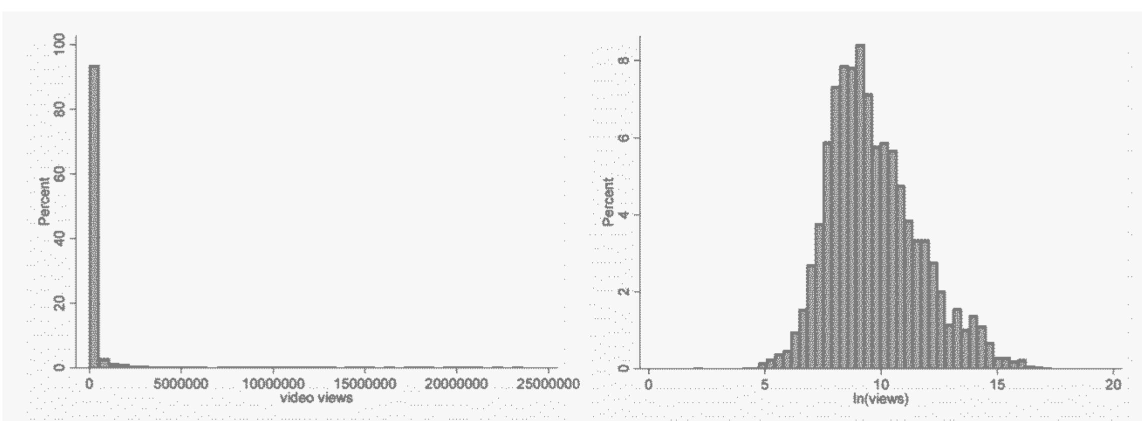

Figure 1.4: Distribution of Views (left) and ln(views) (right) for New Sample Set

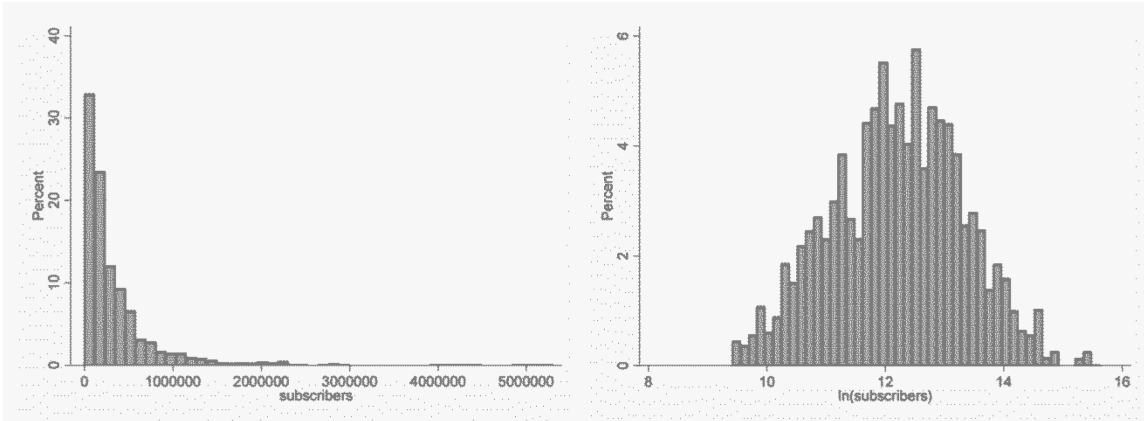

Figure 1.5: Distribution of Subscribers (left) and ln(subscribers) for Top Sample Set

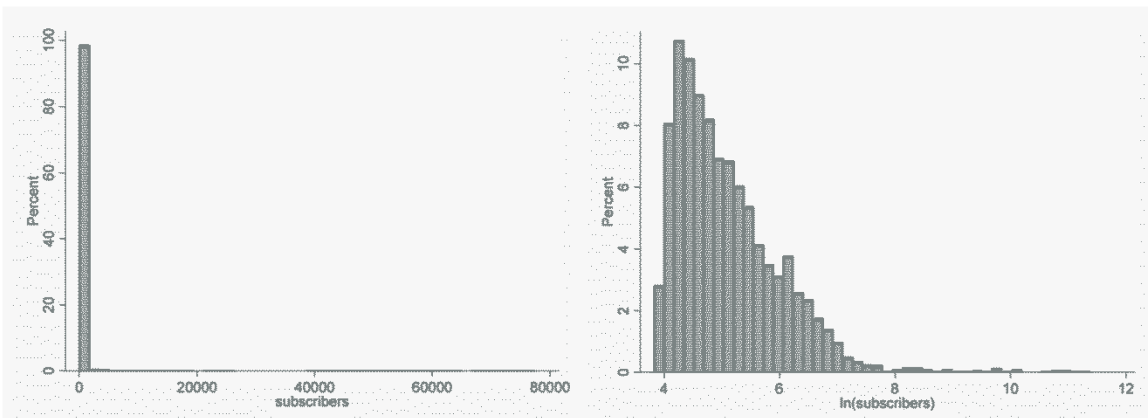

Figure 1.6: Distribution of Subscribers (left) and ln(subscribers) for New Sample Set

Table 1.2: Key Variable Summary

Variable	Mean	SD	Min	Max
v_{it}	0.0102	0.0381	0	4.9457
s_{it}	0.0060	0.0253	-1.1235	3.1761
RV_{it}	6,994	4,115	1	10,000
RS_{it}	5,450	4,524	1	10,000
V_{it}	14.1365	4.5411	0	22.6564
S_{it}	8.8167	3.6701	0	15.4807
P_{it}	0.6153	0.4865	0	1

1.5.3 Censored Data on Reputation

Since vidstatsx.com provides ranking data for only the top 10,000, we can observe the rank in total video views or total subscribers only if it is in the top 10,000. Because we do not observe views rank VR and subscribers rank RS if they are beyond 10,000, we set the unobserved rank to be 10,000. Therefore, the observed-rank data for video views and subscribers RV'_{it} and RS'_{it} are

$$RV'_{it} = \min(10000, RV_{it}) \text{ and } RS'_{it} = \min(10000, RS_{it}). \tag{1.20}$$

To solve this data-censoring problem, we use a Tobit model (Wooldridge 2002) to recover the unobservable RV_{it} and RS_{it} that are beyond 10,000. Our censored Tobit model assumes that RV_{it} given V_{it} and RS_{it} given S_{it} follow normal distributions that

$$RV_{it} = \rho_{v0} + \rho_{v1}V_{it} + z_{it},\ \ z_{it}|V_{it} \sim N(0, \sigma_z^2) \tag{21}$$

$$RS_{it} = \rho_{s0} + \rho_{s1}S_{it} + h_{it},\ \ h_{it}|S_{it} \sim N(0, \sigma_h^2).$$

We use maximum likelihood to estimate ρ_v and ρ_s and recover the value of RV_{it} and RS_{it} before we start estimating our dynamic structural model. The results are reported in Table 1.3. The summary for recovered uncensored values, which are used for structural-model estimation later, is presented in Table 1.4.

Table 1.3: Maximum Likelihood Estimates for Censored Tobi Model

	Parameter estimates (standard errors)
$\boldsymbol{RV_{it}}$	
ρ_{v0} (constant)	35548.48 (71.4041)***
ρ_{v1} (V_{it})	-1850.25 (3.9860)***
Log likelihood	-334687.57
$\boldsymbol{RS_{it}}$	
ρ_{s0} (constant)	46374.72 (264.3245)***
ρ_{s1} (S_{it})	-2244.64 (14.6134)***
Log likelihood	-396394.52

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

Table 1.4: Summary of Uncensored Ranks

Variable	Mean	SD	Min	Max
RV_{it}	13,855	11,042	1	46,374
RS_{it}	7,285	6,532	1	16,222

1.6 Estimation

1.6.1 Two-Stage Estimation Procedure

The two-stage estimation procedure and conditional-choice-probabilities (CCP) based estimator was first proposed by Hotz and Miller (1993). In the two-stage procedure, the first step is to estimate the parameters for choice and transition probabilities. The second step then takes the choices and transition probabilities as givens and estimates the structural parameters in the utility function. Aguirregabiria and Mira (2007), Arcidiacono and Miller (2011), Bajari et al. (2007), and Pakes et al. (2007) developed different estimation methods based on the one proposed by Hotz and Miller (1993). We choose to follow the one suggested by Bajari et al. (2007) because it can deal with continuous-state variables directly without discretization and it is computationally parsimonious compared to other estimation methods.

Let $V_\theta(x_{it}, a_{it})$ denote the choice-specific value function excluding the private shock $\varepsilon_{it}(a_{it})$, which is the expected lifelong utility of choosing a_{it} today and resorting to optimal choices in all future periods. $V_\theta(x_{it}, a_{it})$ can be expressed as

$$V_\theta(x_{it}, a_{it}) = u_{it}(a_{it}, x_{it}) + \beta\mathbf{E}[V_\theta(x_{it+1}, \varepsilon_{it+1})|a_{it}, x_{it}, \varepsilon_{it}(a_{it})]. \tag{1.22}$$

We assume that a contributor's decision is influenced only by that contributor's own state variables, not by the state variables of other contributors. Although interaction and competition generally exist in social media, it is extremely difficult to solve for the equilibrium if interaction and competition are included, given the infinite number of contributors. In this study, we use relative ranking to control for the impact of competition. With these notations, contributor i would optimally choose $a_{it} = 1$ if

$$V_\theta(x_{it}, 1) + \varepsilon_{it}(1) > V_\theta(x_{it}, 0) + \varepsilon_{it}(0). \quad (1.23)$$

We define the policy function (the decision rule for contributors) $\sigma(x, \varepsilon)$ as a mapping from state variables to a binary choice. Since $\varepsilon_{it}(1)$ and $\varepsilon_{it}(0)$ are Type I extreme values, which turns our optimization problem into a logit model, we can recover the choice-specific value functions by inverting the observed conditional choice probabilities at each state (Hotz and Miller 1993) so that

$$V_\theta(x_{it}, 1) - V_\theta(x_{it}, 0) = ln\big(Pr(1|x_{it})\big) - ln\big(Pr(0|x_{it})\big). \quad (1.24)$$

(1.23) is used in the first stage of estimation to derive the optimal decision rule.

Recall that our state variables include video views and subscribers, which gives us a relatively large state space. In this case, a state-by-state inversion approach is likely to generate very noisy estimates of the policy functions. Bajari et al. (2007) suggested that for continuous states, the choice-specific value functions $V_\theta(x_{it}, a_{it})$ can be modeled as flexibly parameterized functions of the action and state variables.

To summarize, the two-stage estimation method can be described as follows. In the first stage, we use fixed effects, OLS, and logit estimations to recover the video contributors' policy functions $\sigma(x, \varepsilon)$, the parameters (γs, σ_ξ^2, σ_ς^2, σ_w^2, and σ_r^2), and the fixed effects μ_i and τ_i. We use the second stage to estimate the structural parameters (αs) that rationalize the contributors' contribution behaviors (δ_1 and δ_2). In this stage, we use forward simulation to derive the minimum-distance estimator that minimizes violations of the optimality conditions (Bajari et al. 2007, Hotz and Miller 199). A single simulated path can be obtained as follows:

1) Starting at observed state $(x_{\mathrm{i}0} = x_{\mathrm{it}})$, draw private shocks $\varepsilon_{it}(a_{it})$ from extreme value distribution for each contributor i.

2) Given the policy function $\sigma(\mathrm{x}, \varepsilon)$ estimated in the first-stage estimation, calculate the optimal action $\mathrm{a}_{\mathrm{it}}^{*}$ given x_{it} and the resulting current-period utility $\mathrm{U}_{\mathrm{it}}(\mathrm{x}_{\mathrm{it}}, \mathrm{a}_{\mathrm{it}}^{*}, \varepsilon_{\mathrm{it}}(\mathrm{a}_{\mathrm{it}}^{*}))$.

3) Draw a vector of error terms for ξ_{it+1}, ς_{it+1}, ω_{it+1}, and r_{it+1} from the corresponding normal distribution, and calculate the resulting new state $x_{\mathrm{i}1}$ using the estimated state-transition parameters.

4) Repeat steps 1–3 for each simulated period.

Subsampling is used to calculate standard errors for structural parameters. Averaging contributor i's discounted sum of utilities over all simulated paths yields an estimate of $V_i(x, \sigma, \theta)$ for any policy function $\sigma(x, \varepsilon)$, including $\sigma^*(x, \varepsilon)$, which is the optimal decision rule derived from first-stage estimation. Because the policy function from the first stage is supposed to be the optimal policy, the following inequality should be satisfied at the true values of structural parameters α_0:

$$g_{\alpha_0}(x_i, \sigma_{\mathrm{i}}^{*}, \alpha) = V_{\alpha_0}(x_i, \sigma_{\mathrm{i}}^{*}, \alpha) - V_{\alpha_0}(x_i, \sigma_{\mathrm{i}}, \alpha) \geq 0. \tag{1.25}$$

The estimator $\hat{\alpha}$ minimizes the objective function below:

$$\hat{\alpha} = \mathrm{argmin}_{\alpha}\left(\frac{1}{N}\sum_{n=1}^{N}(\min\{g_{\alpha_0}(x_i, \sigma_{\mathrm{i}}^{*}, \alpha), 0\})^2\right). \tag{1.26}$$

Although the asymptotic theory requires the simulated time span $T \to \infty$ to drive the simulation error to zero, in practice we can use a finite number of T as long as β^T is relatively small. Given that $\beta = 0.97$, we set $T = 100$ in the simulation.

1.6.2 Identification

Before we estimate the model, we want to provide some intuition for which features of the data allow us to identify particular parameters of the model. The discussion also highlights the assumptions essential for identification. With a panel of observations on new views (v), new subscribers (s), total views (V), total subscribers (S), contribution actions (a), partnership status (P), views rank (RV), and subscribers rank (RS), and the functional form assumptions (Equation (15)) for state transition, the parameters for state transition (γs) can be identified. Once the parameters for state transition (γs) are identified, the error terms (ξ_{it+1}, ς_{it+1}, ω_{it+1}, and r_{it+1}) can be estimated and then their variance (σ_{ξ}^2, σ_{ς}^2, σ_{w}^2, and σ_{r}^2) can be identified.

The common discount factor β cannot be identified because β is highly collinear with the cost parameter α_6. Changing β leads to negligible changes in the value function and estimates of α_6. Both lowering β and raising α_6 can decrease the expected value of contribution and induce the choice of no posting ($a_{it} = 0$). For example, consider a brand new contributor without any views, subscribers, partnership, or reputation as a simple example. For this contributor, the current-period utility for contributing is $-\alpha_6$, and 0 for not contributing. The expected value function for a new state resulting from contributing (denoted as EV1) is likely to be higher than that for a new state resulting from not contributing (denoted as EV0). According to (1.22), the total utility for contributing is $-\alpha_6 + \beta\text{EV1} + \varepsilon_{it}(1)$, and the total utility for not contributing is $\beta\text{EV0} + \varepsilon_{it}(0)$. The difference between the two options is $-\alpha_6 + \beta(\text{EV1} - \text{EV0}) + \varepsilon_{it}(1) - \varepsilon_{it}(0)$. The choice probability is then determined jointly by β and α_6. Therefore, β and α_6 cannot be separately identified from observed contribution behavior. So we take β as known and set it to be 0.97 for forward-looking contributors. We also test the forward-looking assumption against the myopic case where $\beta = 0$ to determine which fits our data better.

Because we do not observe any data on contribution payoffs or costs, we can identify the structural parameters only up to scale. So we normalize the contribution cost α_6 to be 1 and estimate other parameters $(\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5)$ in proportion to α_6. The variation in observed-choice probabilities for contributors with different $v_{it}, s_{it}, P_{it}, RV_{it}, RS_{it}, V_{it}$, and S_{it} allows us to identify $\alpha_1, \alpha_2, \alpha_3, \alpha_4$, and α_5. For example, imagine two contributors with the same state variable except for v_{it}. Because of the way we construct our state-transition process (i.e., v_{it} does not impact state transition), two contributors would form the same expectation for the next-period state variables and thus the same value function for the next period. As β is given, the different choice probabilities are determined by different v_{it}. So α_1 can be identified from the different choice probabilities for the two contributors. Similar interpretations can be drawn for identifying $\alpha_2, \alpha_3, \alpha_4$, and α_5.

1.6.3 Results

Our main results presented in this section are derived on two sample sets pooled together. In the next section for robustness check, we also provide the results from estimation on partner sample and nonpartner sample observations separately. The first step in the two-stage estimation procedure is fixed-effects estimations to derive the parameters for transitions of v, s, OLS regressions for RV and RS, a logit regression for P (if P_{it}=0), and another logit regression to derive conditional-choice probabilities $\mathbf{Pr}(1|x_{\mathrm{it}})$.

Table 1.5 presents results of the fixed-effects estimations. All parameters are significant. Contrary to what we would expect from a Bass diffusion model, V_{it} has a negative impact on v_{it+1}. So we have v_{it+1} decreasing as $\mathrm{V_{it}}$ increases, indicating no herding on total views in choosing contributors. We also find that the number of total subscribers S_{it} has a strong positive impact on new views, indicating herding on total subscribers in choosing contributors. These two findings together suggest total subscribers can be more influential than total video views for a contributor's popularity. Since we use log transformation for these variables, we can interpret the coefficient on $\mathrm{S_{it}}$ as that 1-percent increase in number of total subscribers would lead to a 0.0367-percent increase in number of new views. This impact is slightly stronger when the contributor posts a new video, so that a 1-percent increase in number of total subscribers would lead to a 0.0375 (=0.0372+0.0003)-percent increase in number of new views. While the increase in new views caused by $a_{it}S_{it}$ is statistically significant though, the magnitude of this interaction effect is very small, indicating low loyalty of existing subscribers. We also find significant reputation effects from views rank and subscribers rank, with the effect of views rank slightly stronger than that of subscribers rank.

Similar to the relationship between $\mathrm{V_{it}}$ and $v_{\mathrm{it+1}}$, $s_{\mathrm{it+1}}$ decreases with $\mathrm{S_{it}}$ as well, indicating no herding in viewers' subscription decisions. However, $v_{\mathrm{it+1}}$ has a strong impact on $\mathrm{s_{it+1}}$, suggesting that the more viewers a contributor's videos attract, the more subscribers those viewers can be converted to. The negative coefficient (γ_4) on $\mathrm{v_{it+1}^2}$ further shows that the conversion rate from viewers to subscribers decreases as $\mathrm{v_{it+1}}$ increases. Reputation effect caused by high views rank has a much stronger impact by high subscribers rank. To summarize, these results imply that although viewers herd on contributors with more subscribers when choosing videos to watch, no herding is observed when they make subscription decisions, and that viewers tend to watch videos from and subscribe to providers with high reputations.

Table 1.5: Fixed Effects Estimates for v_{it+1} and s_{it+1}

Transition of v_{it+1}	Parameter estimates (standard errors)
γ_{10} (constant)	0.6379 (0.00361)***
γ_{11} (V_{it})	-0.0314 (0.00027)***
γ_{12} (V_{it}^2)	-0.0023 (0.00001)***
γ_{13} (S_{it})	0.0372 (0.00023)***
γ_{14} ($a_{it}S_{it}$)	0.0003 (0.00003)***
γ_{15} ($1/RV_{it}$)	0.1808 (0.00415)***
γ_{16} ($1/RS_{it}$)	0.1566 (000379)***
σ_ξ^2	0.00064
R^2	37.68%
Transition of s_{it+1}	
γ_{20} (constant)	0.6072 (0.00486)***
γ_{21} (S_{it})	-0.0345 (0.00034)***
γ_{22} (S_{it}^2)	-0.0032 (0.00002)***
γ_{23} (v_{it+1})	0.3180 (0.00413)***
γ_{24} (v_{it+1}^2)	-0.1112 (0.00279)***
γ_{25} ($1/RV_{it}$)	0.3739 (0.00445)***
γ_{26} ($1/RS_{it}$)	0.2238 (0.00371)***
σ_ς^2	0.00046
R^2	32.23%

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

Table 1.6 shows the OLS regression results for $RV, RS,$ and P. Views rank and subscribers rank are both highly correlated with rank in the previous period. New views and new subscribers can improve the contributor's ranking in views and subscribers. On average, a 1-percent increase in v_{it+1} can improve views rank by 2,082, while a 1-percent increase in s_{it+1} can improve subscribers rank by 1,757. For nonpartners, increases in both v_{it+1} and s_{it+1} can increase the chances of becoming partners, but an increase in s_{it+1} can increase the chance by more than the same increase in v_{it+1}.

Table 1.6: OLS Estimates for RV_{it+1} and RS_{it+1} and Logit Estimates for P_{it+1}

Update of RV_{it+1}	Parameter estimates (standard errors)
γ_{30} (constant)	-72.162 (2.5493)***
γ_{31} (RV_{it})	1.0021 (0.00016)***
γ_{32} (v_{it+1})	-2082.3. (54.412)***
σ_w^2	201369
R^2	99.84%
Update of RS_{it+1}	
γ_{40} (constant)	1.9001 (0.21012)***
γ_{41} (RS_{it})	0.9998 (0.00002)***
γ_{42} (s_{it+1})	-1756.8 (5.62460)***
σ_r^2	1471
R^2	99.99%
$\Pr(\mathrm{P_{it+1}} \mid \mathrm{P_{it}} = 0)$	
γ_{50} (constant)	-6.8546 (0.33553)***
γ_{51} (V_{it+1})	0.1280 (0.02845)***
γ_{52} (S_{it+1})	0.2481 (0.05787)***
Log likelihood	-1676

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

For conditional-choice probabilities, we use logit regression of contributors' actions on state variables. The results are presented in Table 1.7. These estimates are used to calculate the empirical probabilities of $\mathbf{Pr}(1|\mathrm{S_{it}})$ and $\mathbf{Pr}(0|\mathrm{S_{it}})$ at each state, which are further used to derive the optimal-decision rules based on (1.23) and (1.24).

In the second stage, we estimate structural parameters in the utility function (αs). The coefficient on contribution cost α_6 is normalized to be 1, and other parameters are estimated relative to α_6. Results are presented in Table 1.8. α_1 and α_2 measure contributors' valuations of new views and new subscribers, respectively. α_3 measures contributors' additional valuation for new views because of revenue sharing. α_4 and α_5 measure contributors' valuations of reputation in video views and subscribers, respectively. We test the forward-looking model, where $\beta = 0.97$ against the myopic model, where $\beta = 0$, and find the data reject the hypothesis that contributors behave as myopic decision makers: The dynamic model with $\beta = 0.97$ produces a statistically significant improvement in ability of the model to fit the data.

Table 1.7: Logit Estimates for Conditional Choice Probabilities

$\Pr(1\|\text{State}_{it})$	Parameter estimates (standard errors)
v_{it}	11.9470 (0.73697)***
s_{it}	5.6445 (0.77343)***
V_{it}	0.0620 (0.04388)
S_{it}	-0.3901 (0.07079)***
RV_{it}	-0.00001 (0.000007)**
RS_{it}	0.00005 (0.00002)***
P_{it}	0.64779 (0.04185)***
μ_i	-1.1107 (0.43982)**
τ_i	4.9745 (0.57017)***
Constant	-0.1034 (0.78716)
Percent correctly predicted	87.21%

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

Table 1.8: Structural Estimates for Utility Function

	Parameter estimates (standard errors)	
	$\beta = 0$	$\beta = 0.97$
α_1 (v_{it})	1.0363 (0.04452)***	1.0075 (0.02661)***
α_2 (s_{it})	1.1110 (0.08394)***	1.0180 (0.06946)***
α_3 $(P_{it}v_{it})$	0.9684 (0.02891)***	1.0125 (0.01663)***
α_4 $(1/RV_{it})$	0.9513 (0.10689)***	1.0212 (0.07315)***
α_5 $(1/RS_{it})$	1.1005 (0.03833)***	1.0321 (0.05096)***
α_6 (a_{it})	1 (normalized)	1 (normalized)
Objective function	52.2342	47.4523
	Myopia test: $\beta = 0$ vs. $\beta = 0.97$	
LR type statistic	9.5638***	

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

According to the structural estimates with $\beta = 0.97$, we can see v_{it}, s_{it}, $P_{it}v_{it}$, RV_{it}, and RS_{it} all have significant impacts on the contributor's utility. Among the two exposure measures, new subscribers can bring in slightly more utility for a contributor than new views. However, this difference is not statistically significant (Table 1.9). If the contributor is able to earn advertising revenue, new views can bring in monetary payment as well as exposure. In this case, new views give the contributor almost twice the utility (1.0075 + 1.0125 = 2.0200) generated by new subscribers (1.0180). Among the two reputation measures, results show a contributor's reputation in terms of video views can bring in slightly more benefits (1.0321) than that in terms of subscribers (1.0212). This difference, however, is not significant (Table 1.9), indicating that reputations in both views and subscribers have a similar importance for contributors.

Table 1.9: Test on Results from $\beta = 0.97$

Conjecture	F statistics
$\alpha_1 = \alpha_2$	0.0197
$\alpha_4 = \alpha_5$	0.0094
Note: Insignificant statistics means that the conjecture is supported.	

*** p = 0.01; ** p = 0.05; * p = 0.10

Compared to these benefits, the relative contribution cost is still high. For example, for a contributor with average v_{it}, s_{it}, RV_{it}, RS_{it} at sample mean (0.0102, 0.0060, 13855, 7285), and no partnership, the average current-period benefit is 0.0166, far from enough for covering contribution cost (fixed as 1). This finding also proves that a forward-looking assumption can better justify most contributors' behaviors than can a myopic one.

1.6.4 Partners vs. Nonpartners

The structural estimates in Table 1.8 are based on the assumption that partners and nonpartners have the same state-transition process. In this section, we test whether the results will still hold if we allow partners and nonpartners to have different state-transition parameters and form different expectations for state variables in the next period. Since we estimate the models for partners and nonpartners separately, we drop the variables for partnership status P_{it} and revenue sharing $P_{it}v_{it}$. The additional valuation for v_{it} as a result of revenue sharing should be reflected as the difference between the coefficients for v_{it} (α_1) for partners and nonpartners. Table 1.10 below presents a comparison of key variables for partners and nonpartners. It shows that although partners have accumulated more total views and subscribers, these nonpartner new providers grow faster with more new views and subscribers.

Table 1.10: Key Variables Summary for Partners vs. Nonpartners

	Nonpartners				Partners			
Variable	Mean	SD	Min	Max	Mean	SD	Min	Max
v_{it}	0.0201	0.0544	0	4.9457	0.0040	0.0203	0	2.4264
s_{it}	0.0114	0.0380	-1.1235	3.1761	0.0026	0.0102	-0.0408	0.4212
RV_{it}*	25082	4706	5901	46375	6837	7502	1	36274
RS_{it}*	14188	1482	2640	16222	2970	4422	1	16184
V_{it}	9.4839	2.0902	0	16.3629	17.0618	2.9257	4.4998	22.6564
S_{it}	4.9770	0.8355	3.8286	10.9436	11.2269	2.5236	3.8501	15.4827

* Recovered from censored data.

Tables 1.11 and 1.12 present results from the first-stage estimation for state transition. Overall, nonpartners and partners experience quite different state-transition processes. For the transition of v_{it+1}, we find that the coefficient estimates for V_{it} and V_{it}^2 for nonpartners have opposite signs from those for partners. Given the range of V_{it} for both sets, $v_{\mathrm{it+1}}$ decreases as $\mathrm{V_{it}}$ increases for partners, whereas it first increases and then decreases as $\mathrm{V_{it}}$ increases for nonpartners. We find similar results on $\mathrm{S_{it}}$ and $\mathrm{S_{it}^2}$ for the transition of $s_{\mathrm{it+1}}$. For nonpartners, views rank has a much stronger positive impact on both new views and new subscribers, while subscribers rank has no significant impact. In fact, we find that subscribers rank has positive impact only on increased views for partners. The explanation for this finding is that as a new contributor gets popular, views rank has a decreasing reputation effect, whereas subscribers rank has an increasing reputation effect. These results overall are consistent with the results for two sets pooling together, although the coefficient magnitudes are quite different between the two sets. Similar conclusions can be drawn for the transitions of RV_{it+1} and RS_{it+1}.

Structural estimates for partners and nonpartners are presented in Table 1.13. The results show the estimate of α_1 to be significantly higher for partners than for nonpartners, confirming that v_{it} generates more benefits for partners because of revenue sharing. For nonpartners, reputation in terms of video views is more important than reputation in terms of subscribers. However, it is the opposite for partners that views rank does not have a significant impact on their utilities but subscribers rank does. This finding suggests that for established top providers, attracting more views can no longer bring in reputational benefits if these views cannot generate new subscribers.

Table 1.11: Fixed Effects Estimates for v_{it+1} and s_{it+1} for Partners vs. Nonpartners

	Parameter estimates (standard errors)	
	Nonpartners	Partners
Transition of v_{it+1}		
γ_{10} (constant)	0.0940 (0.00642)***	0.5515 (0.02588)***
γ_{11} (V_{it})	0.0277 (0.00110)***	-0.0577 (0.00366)***
γ_{12} (V_{it}^2)	-0.0070 (0.00009)***	0.0006 (0.00016)***
γ_{13} (S_{it})	0.0565 (0.00079)***	0.0229 (0.00245)***
γ_{14} $(a_{it}S_{it})$	0.0027 (0.00015)***	0.0001 (0.00002)***
γ_{15} $(1/RV_{it})$	1447.3 (565.30)***	0.0067 (0.00222)***
γ_{16} $(1/RS_{it})$	-320.67 (328.11)	0.0130 (0.00204)***
σ_{ξ}^2	0.001379	0.000186
R^2	31.55%	49.93%
Transition of s_{it+1}		
γ_{20} (constant)	-0.6823 (0.01152)***	0.7816 (0.01230)***
γ_{21} (S_{it})	0.2796 (0.00401)***	-0.1149 (0.00246)***
γ_{22} (S_{it}^2)	-0.0342 (0.00047)***	0.0038 (0.00014)***
γ_{23} (v_{it+1})	0.3643 (0.00715)***	0.2377 (0.00438)***
γ_{24} (v_{it+1}^2)	-0.1259 (0.00497)***	-0.0922 (0.00251)***
γ_{25} $(1/RV_{it})$	5942.4 (354.97)***	0.1109 (0.01734)***
γ_{26} $(1/RS_{it})$	-1147.2 (1312.43)	-0.0189 (0.03148)
σ_{ς}^2	0.001068	0.000064
R^2	32.19%	50.26%

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

Table 1.12: OLS Estimates for RV_{it+1} and RS_{it+1} for Partners vs. Nonpartners

	Parameter estimates (standard errors)	
	Nonpartners	Partners
Update of RV_{it+1}		
γ_{30} (constant)	-4.5366 (0.6212)***	-63.578 (3.3782)***
γ_{31} (RV_{it})	1.0002 (0.00002)***	0.9988 (0.00038)***
γ_{32} (v_{it+1})	-2223.6 (2.4088)***	-1589.0 (146.19)***
σ_w^2	306.02	314837
R^2	99.99%	99.43%
Update of RS_{it+1}		
γ_{40} (constant)	-1.6691 (0.09475)***	2.0665 (0.26953)***
γ_{41} (RS_{it})	1.0001 (0.00001)***	0.9997 (0.00005)***
γ_{42} (s_{it+1})	-1763.8 (0.2462)***	-1692.9 (21.243)***
σ_r^2	2.4963	2304.89
R^2	99.99%	99.99%

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

Table 1.13: Structural Estimates for Partners and Nonpartners

	Parameter estimates (standard errors)	
	Nonpartners	Partners
α_1 (v_{it})	1.0457 (0.01263)***	2.6046 (0.55316)***
α_2 (s_{it})	0.9263 (0.03782)***	1.1378 (0.07739)***
α_4 $(1/RV_{it})$	1.6855 (0.24219)***	1.4318 (0.91940)
α_5 $(1/RS_{it})$	0.9122 (0.09717)***	0.8004 (0.08537)***
α_6 (a_{it})	1 (normalized)	1 (normalized)
β	0.97 (fixed)	0.97 (fixed)

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

1.6.5 Alternative Ranking

In this section, we check the robustness of our main results in Table 8 relative to the alternative-ranking method. We use the ranking among the sample contributors instead of among all YouTube contributors. This test approximates the condition that contributors are compared only to contributors with a similar status. The results from the first-stage estimation for v_{it+1} and s_{it+1} show no significant reputation effects from RV_{it+1} and RS_{it+1}, in terms of attracting viewers and subscribers, so we exclude RV_{it+1} and RS_{it+1} in predicting v_{it+1} and s_{it+1}. Coefficients on other variables are consistent with the main results shown in Tables 1.5 and 1.6. So we present only the structural-estimation results in Table 1.14, which also includes results using rank data among all contributors (Table 1.8) for comparison. According to Table 1.14, results using rank among sample contributors are consistent overall with results using rank among all contributors. Using rank among sample contributors generates higher coefficients for both reputation measures, however, indicating that relative position in comparison to similar contributors may be more important than relative position in comparison to all other contributors in general.

Table 1.14: Structural Estimates Using Different Rankings

	Parameter estimates (standard errors)	
	Rank among sample contributors	Rank among all contributors
α_1 (v_{it})	0.9932 (0.03950)***	1.0075 (0.02661)***
α_2 (s_{it})	1.0132 (0.05752)***	1.0180 (0.06946)***
α_3 $(P_{it}v_{it})$	0.9784 (0.07519)***	1.0125 (0.01663)***
α_4 $(1/RV_{it})$	1.1559 (0.06781)***	1.0212 (0.07315)***
α_5 $(1/RS_{it})$	1.0663 (0.02817)***	1.0321 (0.05096)***
α_6 (a_{it})	1 (normalized)	1 (normalized)
β	0.97 (fixed)	0.97 (fixed)

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

1.6.6 Alternative Reputation Measures

In this section, we check the robustness of our main results in Table 1.8 to alternative reputation measures. We use absolute performance measured by total views and total subscribers, instead of relative performance measured by views rank and subscribers rank, to test how sensitive our estimates are to different reputation measures. The comparison of structural estimation results is presented in Table 1.15. According to Table 1.15, if we use total views and total subscribers as measures for reputation, v_{it}, s_{it} and $P_{it}v_{it}$ no longer have significant impacts on the contributor's utility, suggesting we are not able to disentangle incentives from exposure, revenue sharing, and reputation.

Table 1.15: Structural Estimates Using Different Reputation Measures

	Parameter estimates (standard errors)	
	Reputation measured by views and subscribers	Reputation measured by views rank and subscribers rank
α_1 (v_{it})	0.4880 (1.38873)	1.0075 (0.02661)***
α_2 (s_{it})	3.0034 (12.83798)	1.0180 (0.06946)***
α_3 $(P_{it}v_{it})$	0.2949 (1.36969)	1.0125 (0.01663)***
α_4 $(1/RV_{it})$	45.67 (0.67024)***	1.0212 (0.07315)***
α_5 $(1/RS_{it})$	4.7181 (2.54160)*	1.0321 (0.05096)***
α_6 (a_{it})	1 (normalized)	1 (normalized)
β	0.97 (fixed)	0.97 (fixed)

*** $p = 0.01$; ** $p = 0.05$; * $p = 0.10$

1.7 Discussion and Conclusions

As more and more social media websites introduce revenue sharing as an incentive for content contribution, it is important to understand the effect of this mechanism. It is necessary to recognize that, besides this direct monetary incentive, contributors also receive attention and build reputation in social media. Contributions to and promotions in social media have become successful alternatives to customary career or business paths. For example, traditionally businesses needed to pay for advertising in mass media such as expensive TV commercials. With social media, however, they can create a YouTube channel, Twitter account, or blog to advertise their products or services. The advantages of using social media include lower costs and more frequent interaction with customers. However, how successful businesses or individuals can be using social media depends on whether they can effectively build their reputations. Therefore, it is essential to understand how reputation in social media is measured and how contributors value their reputations.

Based on existing theories of revenue sharing, reputation, and content contribution, we develop a dynamic structural model to identify content contributors' utility function through their content-contribution decisions, using YouTube as the research context. Recovering structural parameters in utility function enables us to study the underlying incentives that fundamentally determine contributors' contribution behaviors. Using a data set from YouTube on 823 top contributors and 1,252 newly registered contributors, we find the following results. First, content contribution in social media is driven by contributors' desires for exposure and reputation. Revenue sharing provides an extra incentive for contributors who join the revenue-sharing program. Second, we prove that content contribution can be better explained by dynamic, forward-looking decision making in which contributors anticipate future benefits as well as immediate rewards. The main advantage of a dynamic structural model is not the estimation of the parameters, but the structural framework for decision making that enables evaluation of counterfactual changes in the environment. Oftentimes, the platform owner is not interested in predicting contribution given current partner program and payment schedule, but interested in what would happen if the program requirements and payment schedule change. Our analysis provides a useful tool to the platform owner for such analyses. Third, it is more appropriate to measure reputation in terms of relative performance than absolute achievement because of the competition effect among all contributors. Comparing with other contributors of similar status is more important than simply comparing with all contributors in general. Fourth, established top contributors enjoy more reputation benefits only if they can attract more subscribers than their competitors. Attracting more viewers can no longer bring in reputation benefits for top contributors if they cannot convert these viewers into subscribers. This result is consistent with the finding that when choosing videos, viewers tend to herd on contributors who have more subscribers rather than those with more views. We also show that no herding exists when viewers choose contributors to subscribe to, suggesting that subscription reflects matching between a contributor’s videos and a subscriber’s preference. Therefore, subscribers rank can better reflect the reputation of a popular contributor than viewers rank.

Our results have important implications for practitioners and researchers in social media. First, revenue sharing and reputation can both be used to motivate content contribution, especially for video-content contribution, which is costly for contributors. For popular contributors, reputation benefits are brought in by subscribers rank but not by viewers rank. Without revenue sharing, these contributors would be more interested in building loyal subscribers and reputation in terms of subscribers than attracting general viewers. However, this may not be consistent with the platform owner’s interest to generate more engagement for advertising. Revenue sharing can motivate contributors to attract general viewers and to cater to advertisers’ interests. For websites that do not offer revenue sharing, the capability to provide a powerful platform for contributors could still attract contributors pursuing recognition and reputation. Besides revenue sharing, YouTube has so many contributors also because it is the most influential video-sharing website in the world. The Huffington Post and Twitter are not paying any contributors but still have large contributor bases. Second, motivating contributors with potential future rewards is important. In fact, we have already seen such examples in practice. In 2010, YouTube had a $5-million-grant program for "new and emerging YouTube partners." It also holds various contests from time to time. Break.com has the policy that contributors get paid if their videos are posted on the front page. Such possible rewards to come are necessary for mitigating high contribution costs today, which can induce high-quality content. Third, websites can emphasize contributors' reputations by comparing contributors to each other. More importantly, websites can identify groups of similar contributors with comparable popularities and content interests to provide ranking for different groups. Such specific comparison is more useful for encouraging content contribution than general comparison. In addition, social-media websites should give more weight to subscriptions than to views so they can make better recommendations.

1.8 Future Research

The analysis carried out in this research can be extended along a number of dimensions. First, we consider only the aggregate viewership for all contributors' videos. Distribution of views across videos may also have an impact on the contributor's utility though. For instance, a single hit video might be more valuable than several average ones. Data at the video level are necessary to carry out a more detailed analysis. Second, we simplify the contribution decision to be a binary choice between contributing or not, while in reality the decision is much more complex. A contributor needs to choose the quantity and quality of the content to contribute as well. Another simplification is that, due to data availability, we only consider the contributor's decision on posting videos but not producing videos. In reality, a contributor makes two sequential decisions, first determine whether it will be worthwhile to produce a video for potential distribution through social media sites and second determine when will be the optimal time to post the video once the video is produced. The two sequential decisions are clearly dependent on each other. It will be worthwhile for future research to analyze the two decisions jointly. Third, we measure reputation using both ranks among all contributors and ranks among contributors with similar popularities in this research. It could be more meaningful to use ranks among contributors working on similar topics. More importantly, our operationalization of reputation using rank measures is a relatively crude approach. Reputation, by its nature, is latent and contains multiple dimensions. Incorporating other dimensions of reputation such as viewer ratings and related WOM measures could help improve the model. While our preliminary analysis using a subset of the data with video rating data indicates that rating has no significant impact on viewership or subscription, a better understanding and modeling of reputation in online social media will be an interesting topic for future search. Finally, we do not look into the details of viewers' choices of videos and channels, which could shed light on the dynamics of content contribution and competition between channels. These limitations provide exciting opportunities for future research.

Chapter 2: Distinguishing Social Learning from Network Effects in Social Media

2.1 Introduction

With new products such as consumer goods, food, pharmaceutical goods, financial services, and movies constantly flooding the markets, consumers face an already overwhelmingly large and rapidly growing choice set. Meanwhile, with the prolific use of social media, consumers obtain information about products from social sources in the forms of product reviews and friends' recommendations. Therefore, marketers tend to use these forms of social contagion to influence consumers' perception and behavior.

Several studies have shown the presence of social contagion in new product adoption (e.g., Godes and Mayzlin 2009; Goldenberg et al. 2009; Iyengar, Van den Bulte, and Valente 2011). We aim to take a step further to test different mechanisms of social contagion: How individuals' decisions are affected by peers? In our context, social contagion happens mainly through two channels: (1) Social learning, the process in which consumers obtain knowledge about a product's quality through peers; (2) Network effect, the phenomenon that the value of a product increases as the number of its users increases. Which mechanism exists or dominates depends on the specific product in question. When choosing a mobile network operator, network effects may dominate because of free mobile-to-mobile calling. When purchasing an HDTV, social learning becomes the primary force because consumers are mainly concerned about the quality. These two mechanisms have different implications for marketing strategies: For products with strong network effects, creating a large user base is crucial in attracting new adopters, while for products with prevailing social learning, generating positive word-of-mouth (WOM) is the key.

In this study, we differentiate between social learning and network effects in the context of YouTube, the largest online video sharing website. Choosing online videos to watch is one of the most common choices viewers make every day. Given the vast reservoir of online videos, how to choose videos to watch can become a complicated issue. On the one hand, consumers receive various information from friends and infer video quality through social learning. On the other hand, frequent social sharing creates direct or indirect network effects where a video becomes a fad. For example, when a video goes viral, users have strong incentives to watch it so they have something to discuss in social encounters.

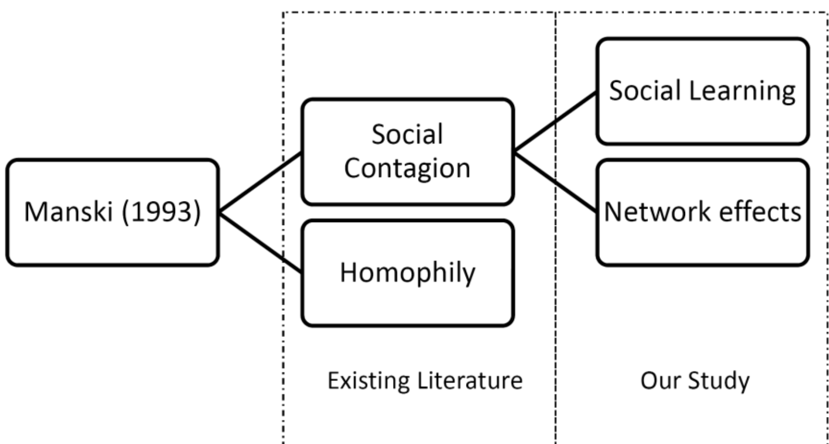

Figure 2.1: The Conceptual Framework of Social Learning and Network Effects

Figure 1 shows the conceptual framework of our study. Most existing studies on social contagion focus on the Manski problem (Manski 1993): distinguishing general social contagion from homophily - the tendency of individuals to associate with similar others (Aral et al. 2009, Aral and Walker 2011). Few of them differentiate between the two mechanisms of social contagion: social learning and network effects. Social learning affects consumers through the quality information conveyed by peers, whereas the network effects influence consumers according to the size of the user base. Although these mechanisms lead to similar empirical outcome, their implications are vastly different. If social contagion is generated mainly by network effects, then seeding strategies, which determine the initial set of targeted consumers, would by implication have a strong influence on the success of viral marketing. Ho et al. (2012) show that a firm can amplify social contagion and accelerate product purchases by offering introductory discounts. However, if social learning is the dominant effect, seeding would not be effective unless the initial consumers generate positive word-of-mouth. Qiu and Whinston (2012) demonstrate that consumers can infer that the high demand of their peers is caused by the introductory discount rather than the high product quality. Both cases are theoretically plausible and need to be empirically distinguished.

To the best of our knowledge, our study is the first to disentangle social learning and network effects in the context of online video sharing. Because of the lack of pre-release marketing effort, these two types of social contagion are particularly important for user-generated content (UGC). Our empirical results suggest that both mechanisms through which social contagion works are important. We further find that social learning is more pronounced when consumers are less certain about video quality, and videos with attention grabbing content bring on higher network effects. Eliaz and Spiegler (2011) examine the incentives of a firm to offer low quality products as "pure attention grabbers." Our study confirms their theoretical results: attention grabbers provoke discussions and go viral through network effects. The implications derived from studying YouTube can carry over to other consumer choice problems as well.

The rest of the chapter is organized as follows. We review related literature in Section 2.2. In Section 2.3, we outline an analytical model that motivates our empirical hypotheses. Section 2.4 describes the YouTube data. In Section 2.5, we present the identification strategy and empirical results. Some applications are explored in Section 2.6. Section 2.7 concludes.

2.2 Literature Review

Manski (1993) discuss an econometric challenge of identifying social contagion: Is a person's behavior caused by his social reference group, or does it simply reflect the same movement in his reference group? The observation that individuals belonging to the same group tend to behave similarly might result from social contagion, exogenous contextual effects, or homophily[3]. Failure to account for contextual effects or homophily might lead to an overestimation of the effect of social contagion.

These confounding effects are difficult to distinguish, and the identification of social contagion often requires strong parametric assumption or rich data collection. Aral, Muchnik, and Sundararajan (2009) distinguish influence-based contagion from homophily-driven diffusion using a dynamic matched sample of global instant messaging users. Iyengar, Van den Bulte, and Valente (2011) distinguish social contagion from homophily and exogenous contextual effects in prescribing behavior among networks of doctors.

Within the framework of social contagion, studies have been focusing on distinguishing social learning from other contagion mechanisms such as saliency effect

[3] Among these three effects, only social contagion can generate "social multiplier" with a positive feedback (Manski 1993).

(i.e., observed choices are more salient than alternative choices), conformity concerns (i.e., the social pressure to adopt the choice made by the majority), and network effects.[4] Although some studies provide evidence of the existence of social contagion in the diffusion of UGC, none of them look into the two specific contagion mechanisms, social learning and network effects, each of which may have different managerial implications.[5] Using a diffusion model for YouTube videos, Susarla et al. (2011) demonstrate that social networks affect economic outcomes by structuring the information available to other users, which influences their decisions, perceptions, and behaviors. Goldenberg et al. (2012) show that the stream of people's chatter from social broadcasting networks facilitates social learning among a much broader peer group than has traditionally been possible. Liu-Thompkins and Rogerson (2012) identify three factors that affect the diffusion outcomes of YouTube videos: network structure, content characteristics, and author characteristics.

For UGC, understanding whether the popularity of the content makes it valuable (network effects) or the value of the content makes it popular (social learning) is pivotal. By distinguishing between network effects and social learning, our study contributes to the understanding of different social contagion mechanisms in the diffusion of UGC.

---

[4] Cai, et al. (2009) use a field experiment to distinguish observational learning from saliency effect. Van den Bulte and Stremersch (2004) study different social contagion mechanisms using a meta-analysis of publications on new product diffusion and find evidence for status concerns and social-normative pressures but not for social learning under uncertainty. Iyengar, Van den Bulte, and Choi (2012) differentiated between social learning and normative influence in the adoption of a new drug.

[5] Online WOM, especially online user reviews, has become an important channel of social learning for consumers, which directly affects their behavior. Duan et al. (2009) characterize the WOM process of the movie industry through a dynamic simultaneous equation system. Using online book reviews from Amazon.com and Barnesandnoble.com, Chevalier and Mayzlin (2006) find that an improvement in reviews for a book at one site leads to a relative increase in its sale at that site. Goldenberg, Libai, and Muller (2010) separate network effects from WOM learning using empirical evidence in the growth of fax machines, CB radios, CD players, DVD players, and cellular service. Moretti (2011) shows that social learning is a more important determinant of sales in the movie industry than network effects.

2.3 A Theoretical Framework of Social Learning and Network Effects on YouTube

A video can go viral either because of social learning or network effects. This section examines different implications of the two mechanisms.

2.3.1 A Model of Social Learning on YouTube

YouTube videos are experience goods whose quality cannot be fully observed by consumers ex ante, but can be ascertained upon consumption. Therefore, before consumption, consumers are never completely sure about the quality, but they can always acquire useful information from friends who have already watched the videos. Banerjee (1992) examines the social learning that occurs through observing other people's behaviors. Our approach is somewhat different. We capture the learning process with a Bayesian learning model, where each consumer receives feedback from peers and updates the prior belief of the video quality.[6] We extend the observational learning literature by adding underlying social networks. People make inferences about the quality of a video according to the information within their social networks. Because consumers on YouTube face a large and growing choice set, we assume that they have limited information about a video's existence.[7] The probability that a consumer watches a video is the product of two probabilities: the probability that he is aware of the video and the probability of watching the video conditional on being aware of the video. Figure 2.2 shows the timeline. We focus on the information updating process at time 1 and time 2. The process proceeds in the same way at time 3, 4, …, T.

---

[6] Following Banerjee (1992) and Moretti (2011), the timing of consumption is exogenously given, and we do not consider the strategically behavior of delaying the decision making process to obtain more feedback.

[7] As of August 2012, on average, about 72 hours of video are uploaded to YouTube every minute, and the number is still growing, see http://www.youtube.com/t/press_statistics.

We first describe the decision process of conditioning on that they are aware of the existence of the video. The utility that a representative consumer i obtains from watching some YouTube video j is

$$u_{ij} = V_j + \eta_{ij}, \eta_{ij} \sim N\big(0,1/\rho_\eta\big),$$

where V_j is the latent quality of the video, and η_{ij} represents the unobserved taste heterogeneity. This video is published on YouTube at time 1. Individuals share a common prior on the quality of the video, given by

$$V_j \sim N\left(X_j'\beta, 1/\rho_{V_j}\right),$$

where X_j is a vector of the observable characteristics of video j before watching. $X_j'\beta$ is the ex-ante expectation of quality, and ρ_{V_j} is the precision of prior for video j.

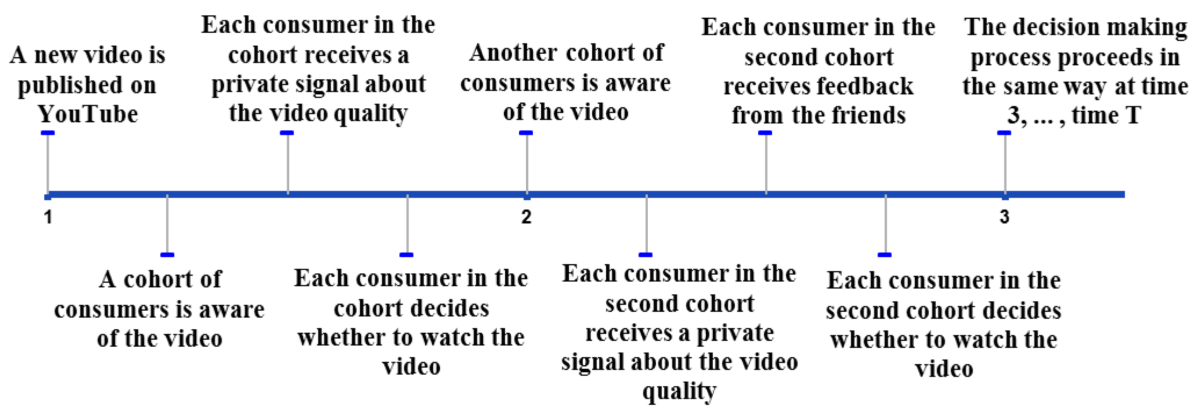

Figure 2.2: Timeline

Before making a decision, each consumer observes a conditionally independent private signal of the quality:

$$S_{ij} = V_j + \varepsilon_{ij},\ \varepsilon_{ij} \sim N(0,1/\rho_{\varepsilon_j}), \tag{3.1}$$

where ρ_{ε_j} is the precision of consumer i's information source for video j. The signal errors ε_{ij} are independent across consumers. Consumers update the prior according to Bayes' rule. $E_t[u_{ij}|I_t]$ represents consumer i's expected utility of video j at time t given the information set at time t, I_t. Notice that the video is newly published, so no social learning occurs at time 1. Conditional on being aware of video j, a consumer chooses to watch it if the ex-ante expected utility is no less than the opportunity cost of watching video j, c_{ij}. Therefore, if a consumer is aware of video j at time 1, then he watches video j if

$$E_1[u_{ij}|I_1] = \frac{\rho_{V_j}}{\rho_{V_j} + \rho_{\varepsilon_j}} X_j'\beta + \frac{\rho_{\varepsilon_j}}{\rho_{V_j} + \rho_{\varepsilon_j}} S_{ij} \geq c_{ij}.$$

Accordingly, the probability that a consumer watches video j at time 1 is:

$$\Pi_1 = \Pr\left(E_1[u_{ij}|I_1] \geq c_{ij}\right) = \Phi\left(\frac{\frac{\rho_{V_j}}{\rho_{V_j} + \rho_{\varepsilon_j}} X_j'\beta + \frac{\rho_{\varepsilon_j}}{\rho_{V_j} + \rho_{\varepsilon_j}} V_j - c_{ij}}{\sqrt{\rho_{\varepsilon_j} / \left(\rho_{V_j} + \rho_{\varepsilon_j}\right)^2}}\right),$$

where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal distribution.

Then we need to model the probability that a consumer is aware of video j at time 1, denoted as W_{j1}. W_{j1} is a function of the characteristics of video j at time 1, such as the video ratings, the number of YouTube Favorites, and the number of video comments. The proportion of informed consumers for video j at time 1 is given by: $P_{j1} = W_{j1}$. Given a large number of viewers on YouTube[8], the number of views of video j at time 1 is

$$views_{j1} = P_{j1}\Pi_1 N = W_{j1}\Pi_1 N,$$

where N is the number of potential consumers.

With social learning, consumers have more information at time 2 because they receive feedback from friends. The underlying social network $\Gamma = (M, L)$ is given by a finite set of nodes $M = \{1,2,\ldots,N\}$ and a set of links $L \subseteq M \times M$. Each node represents a

[8] Over 800 million unique users visit YouTube each month, see http://www.youtube.com/t/press_statistics. According to the law of large numbers, we can calculate the number of video views. If the number of potential consumers is not sufficiently large, Chebyshev's inequality can give us a bound of views.

consumer. The social connections between the consumers are described by an $N \times N$ matrix denoted by $g \in \{0,1\}^{N \times N}$ such that:

$$g_{ij} = \begin{cases} 1, & \text{if } (i,j) \in L \\ 0, & \text{otherwise} \end{cases}.$$

Let $M_i(g) = \{j \in N: g_{ij} = 1\}$ represent the set of friends of consumer i. We assume that consumer i has k friends, where $k = \#N_i(g)$. Among them, k_1 friends have watched the video at time 1. These friends communicate to consumer i their ex-post utilities after watching the video, u_{mj}, where m =1, 2, 3, ... ,k_1. r_1 friends were aware of the video, but decided not to watch the video at time 1, and they are indexed by m= $k_1 + 1, k_1 + 2$, ... , $k_1 + r_1$. The friends that have decided not to watch the video also provide valuable information: their expected utilities are less than the opportunity cost of watching the video. The number of friends who were unaware of the video at time 1 is $k - k_1 - r_1$.

As a result, at time 2, consumer i's information set consists of the ex-post utilities of some friends, the number of friends who decided not to watch the video, and the number of friends who were unaware of the video. Combining these three pieces of information, consumer i estimates the quality by maximizing the likelihood of the observed evidence:

$$LH\left[u_{mj}, m = 1,2,3, \dots, k_1;\ r_1\ ; k - k_1 - r_1 \middle| V_j\right]$$

$$= \Pi_{m=1}^{k_1} f\left(u_{mj}\right) \cdot \Pi_{m=k_1+1}^{k_1+r_1} \Pr\left(E_1\left[u_{mj}\middle|I_1\right] < c_{ij}\right) \cdot \Pi_{m=1}^{k_1+r_1} P_{j1} \cdot \Pi_{m=k_1+r_1+1}^{k} (1 - P_{j1}),$$

where $f\left(u_{mj}\right)$ is the likelihood of observing u_{mj}. The maximum likelihood estimator, G_{ij2}, is an estimate of V_j. It is unbiased and asymptotically normal:

$$G_{ij2} \sim N\left(V_j, 1/d_{i2}\right),$$

where $d_{i2} = -E\left[\left(\frac{\partial \ln L}{\partial V_j}\right)^2\right]$ (Amemiya 1973).

If a consumer is aware of video j at time 2, his expected utility becomes the weighted average of the prior mean, his private signal, and the maximum likelihood estimator:

$$E_2[u_{ij}|I_2] = \frac{\rho_{V_j}}{\rho_{V_j} + \rho_{\varepsilon_j} + d_{i2}} X_j'\beta + \frac{\rho_{\varepsilon_j}}{\rho_{V_j} + \rho_{\varepsilon_j} + d_{i2}} S_{ij} + \frac{d_{i2}}{\rho_{V_j} + \rho_{\varepsilon_j} + d_{i2}} G_{ij2}.$$

Note that as time goes on, consumers put less weight on the prior mean. Because consumers receive more information at time 2, the prior becomes a less important factor in the decision making process. The probability that consumer i watches video j at time 2 is:

$$\Pi_2 = \Pr\left(E_2[u_{ij}|I_2] > c_{ij}\right) = \Phi\left(\frac{\frac{\rho_{V_j}}{\rho_{V_j} + \rho_{\varepsilon_j} + d_{i2}} X_j'\beta + \frac{\rho_{\varepsilon_j} + d_{i2}}{\rho_{V_j} + \rho_{\varepsilon_j} + d_{i2}} V_j - c_{ij}}{\sqrt{\left(\rho_{\varepsilon_j} + d_{i2}\right) / \left(\rho_{V_j} + \rho_{\varepsilon_j} + d_{i2}\right)^2}}\right).$$

If a consumer is aware of video j at time T, the decision making process proceeds in the same way. Consumer i has k_t friends who decide to watch the video at time t, r_t friends who decide not to watch the video at time t, and $k - k_t - r_t$ friends who are unaware of the video, where $t = 1, 2, 3, \ldots, T-1$. The probability that consumer i watches video j at time T is:

$$\Pi_T = \Pr\left(E_T[u_{ij}|I_T] > c_{ij}\right) = \Phi\left[\frac{\alpha_T X_j'\beta + (1-\alpha_T)V_j - c_{ij}}{g_T}\right].$$

where $g_T = \sqrt{\left(\rho_{\varepsilon_j} + \sum_{t=2}^{T} d_{it}\right) / \left(\rho_{V_j} + \rho_{\varepsilon_j} + \sum_{t=2}^{T} d_{it}\right)^2}$, and $\alpha_T = \frac{\rho_{V_j}}{\rho_{V_j} + \rho_{\varepsilon_j} + \sum_{t=2}^{T} d_{it}}$. The proportion of informed consumers at time T is given by: $P_{jT} = P_{jT-1} + (1 - P_{jT-1})W_{jt}$. The number of views of video j at time T is:

$$views_{jT} = (P_{jT} - P_{jT-1})\Pi_T N, \tag{2.2}$$

where $P_{j0} = 0$.

A few remarks need to be made here. In the process of social learning, α_T is the weight that consumers put on the ex-ante prior. It is evident that α_T decreases with T. As time T grows, the probability of watching videos relies less on the ex-ante prior and more on social learning. If the revealed quality of the video is higher than the mean of the ex-ante prior, $V_j > X_j'\beta$, we call it a positive surprise. If the revealed quality of the video is lower than the prior mean, $V_j < X_j'\beta$, we call it a negative surprise. Social learning is a process of adjusting beliefs about the quality. Thus, we have the following proposition:

Proposition 2.1. *In the presence of social learning, if a positive surprise is sufficiently large ($V_j \gg X_j'\beta$), then Π_T is increasing in T. If a negative surprise is sufficiently large ($X_j'\beta \gg V_j$), then Π_T is decreasing in T.*

The intuition is straightforward. In our model, consumers learn about the surprise over time, and a positive surprise increases the expected quality of the video as time goes by. Therefore, the probability of watching the video increases. Similarly, a negative surprise reduces the expected quality over time, and the probability of watching the video decreases.

Other things being equal, we consider some video j with a large positive surprise and some video j' with a large negative surprise. According to (2.2), we find that:

$$\ln views_{jT} - \ln views_{jT-1} > \ln views_{j'T} - \ln views_{j'T-1}.$$ [9]

Therefore, we have the following testable hypothesis from the theoretical prediction:

[9] The logarithm growth rates are widely used in economic modeling and empirical study. They are good approximations for percentage growth rates.

Hypothesis 2.1. *In the presence of social learning, the growth rate of views of a video that has a positive surprise is higher than the growth rate of views of a video that has a negative surprise.*

In our model of social learning, we can also consider the impact of the consumers' prior. The intuition is that social learning is more important when consumers have more diffuse priors.

Proposition 2.2. *(a) If the positive surprise is sufficiently large,* $\Pi_{T+1} - \Pi_T$*, is decreasing in the precision of the prior,* ρ_{V_j}*. (b) Similarly, if the negative surprise is sufficiently large,* $-(\Pi_{T+1} - \Pi_T)$ *is decreasing in* ρ_{V_j}*.*

The incremental probability, $\Pi_{T+1} - \Pi_T$, can measure the effect of social learning. This proposition shows that the effect of social learning is more pronounced for videos with less precise priors. If a consumer is very uncertain about the quality of a video, the value of social learning is large: The additional information he learns from his friends should be more valuable than the case when he knows the quality precisely. An increase in the precision of the prior makes the additional information from friends less valuable. Thus, social learning should be more valuable among videos that are less familiar to consumers, and we have the following empirically testable hypothesis:

Hypothesis 2.2. *In the presence of social learning, the positive surprise has a greater impact on videos with less precise priors.*

2.3.2 A Model of Network Effects on YouTube

Social contagion can be driven by either social learning or network effects. Network effects in our context mean that the utility of watching a video directly depends on the number of consumers who have watched the video, irrespective of their reasons for the choice of watching the video. Both social learning and network effects can be recognized as a form of causal social contagion. However, their underlying mechanisms are different. The essence of network effects is payoff externality, which implies that the value of the service depends directly on the consumption choices made by some other consumers. For example, a consumer might enjoy discussing a video with his peers. In this case, the actions of other consumers do not convey any quality information about the video. Social learning is a different form of causal social contagion, which focuses more on information externality instead of payoff externality. In the presence of social learning, consumers make inferences about the quality of a video by observing other people's choices and comments. They care about the actions of others only because they can receive information about the quality from the peers.

We modify the theoretical model of social learning to introduce network effects and assume that the utility consumer i obtains from watching YouTube video j at time T is given by:

$$u_{ijT} = V_j + \eta_{ij} + \delta \sum_{t=1}^{T-1} (P_{jt} - P_{jt-1}) \Pi_t N \, , \; \eta_{ij} \sim N\left(0,1/\rho_\eta\right),$$

where $(P_{jt} - P_{jt-1}) \Pi_t N$ is the number of consumers who have watched the video at time t. The consumer derives utility from the total number of consumers who have watched the video. The parameter δ measures the impact of network effects. If $\delta > 0$, then network effects exist. If $\delta = 0$, then the impact of network effects is insignificant. For pure network effects model, we assume no social learning, and, consequently, consumers do not receive feedback from peers.

Under network effects, the probability that a consumer watches video j at time t is given by:

$$\Pi_T = \Phi\left(\frac{\frac{\rho_{V_j}}{\rho_{V_j}+\rho_{\varepsilon_j}}X_j'\beta + \frac{\rho_{\varepsilon_j}}{\rho_{V_j}+\rho_{\varepsilon_j}}V_j + \delta\sum_{t=1}^{T-1}(P_{jt}-P_{jt-1})\Pi_t N - c_{ij}}{\sqrt{\rho_{\varepsilon_j}/\left(\rho_{V_j}+\rho_{\varepsilon_j}\right)^2}}\right)$$

It is evident that Π_T is increasing in T under network effects no matter what the values of $X_j'\beta$ and V_j are.

One way to empirically identify network effects is to examine sequential correlation in video views such that the views at time t is positively correlated with the lagged cumulative video view, $(P_{jt} - P_{jt-1})\Pi_t N$. However, we would not be able to uniquely identify network effects without controlling for social learning. To achieve this goal, we perform a two-stage test. In the first stage, we test whether social learning exists. Recall that in Hypothesis 2.1, with social learning, a video with a positive surprise has a higher growth rate than a video with a negative surprise. However, under network effects, a significant difference in the growth rates resulting from surprises would be absent. Therefore, if the empirical evidence supports Hypothesis 2.1, then it would show the existence of social learning. If Hypothesis 2.2 is also confirmed, it would provide additional evidence for social learning.

In the second stage, we test whether network effects exist using instrumental variable as a source of exogenous variation for existing levels of views. If social contagion is purely driven by social learning, the growth rate of video views should remain unchanged when the surprise does not reflect information about video quality. It is because consumers learn nothing from the surprise that does not contain any quality information. However, in the presence of network effects, a negative surprise changes the viewer base (the number of consumers who have already watched the video), and thus lead to a lower growth rate of video views. Through examining Hypothesis 2.3, we can provide the evidence of the existence of network effects. We will describe the two-stage test in more detail in Section 2.5.

Hypothesis 2.3. *In the presence of network effects, the negative surprise lowers the growth rate of video views even if the negative surprise does not reflect information about video quality.*

2.4 Data

To empirically test the theoretical model, we focus on new videos that were posted during our data collection period. As the world's largest video viewing and sharing website, YouTube has enormous numbers of videos, which makes it infeasible to select a random sample set of videos. Instead, we focus on the most active providers by selecting the top 1,000 YouTube providers (in terms of total video views) identified for June 2011.[10] We collect a daily panel of data on these providers for one month, from March 1, 2012 to March 31, 2012. Our sample includes a total of 302 new videos published on March 1, 2012, by these top providers. We use one day as the time unit of analysis to capture the fast-changing nature of online videos.

The provider level data include provider ID, data collection date, date when the provider joined YouTube, number of subscribers to the provider's channel, total views of all the provider's videos, total views of the provider's channel page, number of videos, number of friends, number of subscriptions the provider has to other providers, channel views rank, and video views rank. The video level data include video ID, data collection date, date when the video is posted, the provider of the video, number of views, category

[10] In this study, we focus on social learning and network effects given that consumers are aware of the video. We do not study how consumers become aware of a video. That is why we select the videos published by the top 1,000 YouTube providers as our sample.

in which the video belongs, video length, whether the video has an in-stream ad, average rating, number of times the video is favorited by viewers, and number of comments. All videos in our sample are published on March 1, 2012. Because YouTube Analytics data is updated daily, the first day in our analysis is March 2, 2012. Summary statistics of the video characteristics at the beginning of our data collection period are reported in Table 2.1. We assume that each viewer watches a video only once. Consumers may repeatedly watch a video. However, Susarla et al. (2011) argue that the bias caused by repeated viewings is small by taking logs of views. Table 2.2 provides summary statistics of the characteristics of our YouTube providers.

Table 2.1: The First-Day Video Characteristics

Variables	Mean	Std. Dev.	Min	Max
Number of video views	2,497.20	8,524.969	2	107,628
Video rating	4.692	0.5066	1.76	5
Number of times the video being favorited	139.507	495.910	0	6,800
Number of comments	384.173	1030.539	0	9,832
In-stream ads (yes-1, no-0)	0.5359	0.4995	0	1

Table 2.2. The First-Day Chanel Characteristics

Variables	Mean	Std. Dev.	Min	Max
Total views of the provider's channel page	1.32e+07	2.01e+07	3,175,291	1.75e+08
Total views of all the provider's videos	1.87e+08	2.38e+08	3,690,640	1.55e+09
Number of subscribers to the provider's channel	298,743.8	459,699.2	9,200	5,109,145
Number of subscriptions by the provider to others	183.6144	1,063.834	0	17,641
Number of videos	267.9837	289.788	1	969
Number of friends	16,976.32	22,916.99	0	120,570

2.5 Empirical Framework

2.5.1 Identification of the Surprise

In our theoretical model, the surprise is defined as the difference between the revealed quality and the prior mean, $V_j - X_j'\beta$. Following Moretti (2011), we empirically define the surprise as the difference between realized video views and predicted video views at time 1 (March 2, 2012).[11] In our study, we obtain a measure of the surprise using the prediction errors. More specifically, we use the characteristics of YouTube providers (channel) to predict the video views on the first day. Then, we obtain the residual from a regression of first-day log views on the characteristics of YouTube providers at time 1 as a measure of video-specific surprise. The residual is considered as the difference between the realized video views and the predicted video views at time 1 and thus the measure of surprise. However, the residual may change with different functional forms used for the predicted views. Therefore, as a robustness check, we test our hypotheses with different regression forms for video views.

The characteristics of YouTube providers are reasonable measures of expected video quality. Most consumers on YouTube subscribe to some channels they like. By subscribing to a channel, they receive updates and are informed when new videos are uploaded by the provider. Consumers expect that a high-quality channel will upload high-quality videos.

Table 2.3 illustrates the first-stage regression results of first-day (March 2, 2012) log video views on channel characteristics at time 1. The residual obtained from this first-stage regression is used in the tests for Hypothesis 2.1. The channel characteristics include the log of total views of channel j's videos, $lvviews$; the log of total views of the

[11] In finance literature, Jegadeesh and Titman (1993) document a momentum phenomenon: Firms reporting positive earnings surprises outperform firms reporting negative earnings surprises.

provider's channel page, $lcviews$; the log number of uploaded videos of the channel, $lvideos$; the number of the provider's subscribers, $subs$; and the number of other providers the provider subscribe to, $subscriptions$. Column 1 in Table 2.3 shows the regression results. Column 2 and 3 indicate that the results are robust to other regression specifications.

Table 2.3: Identification of the Surprise: First-Stage Regression

	(1)	(2)	(3)	(4)
lvviews	0.444***	0.445***	0.449***	0.0952***
	[12.83]	[12.89]	[13.0]	[3.112]
lcviews	0.297***	0.297***	0.291***	0.0799*
	[4.809]	[4.798]	[4.710]	[1.661]
lvideos	0.0175	0.0176	0.0179	0.0171
	[0.564]	[0.568]	[0.579]	[1.512]
subs	3.26e-08	3.29e-08		4.43e-08
	[0.138]	[0.140]		[0.530]
subscriptions	2.71e-06			1.33e-05
	[0.0360]			[0.504]
constant	1.630	1.629	1.410	5.090***
	[0.684]	[0.685]	[0.787]	[5.886]
Observations	302	302	302	302
R-squared	0.545	0.545	0.544	0.528

t-statistics in brackets: *** *p<0.01,* ** *p<0.05,* * *p<0.1*

One may wonder if the residual is a robust measure of surprise. We also use another empirical definition of surprise to show the robustness: the difference between realized video ratings and predicted video ratings at time 1. Similarly, we run a regression of first-day video ratings on the characteristics of YouTube providers. This regression is shown in Column 4 in Table 2.3. Then, we obtain the residual from it as a measure of surprise. In Section 2.5.2 and Section 2.5.3, we will show that our empirical results are robust to different empirical definitions of surprise. One may also wonder if the residual could contain a number of omitted variables that are observable to the viewer but not measured by the researcher. If it is the case, the residual will be systematically correlated with viewership. We run a regression of log views on the residual at time 1, and the coefficient is not significant (p value = 0.331).

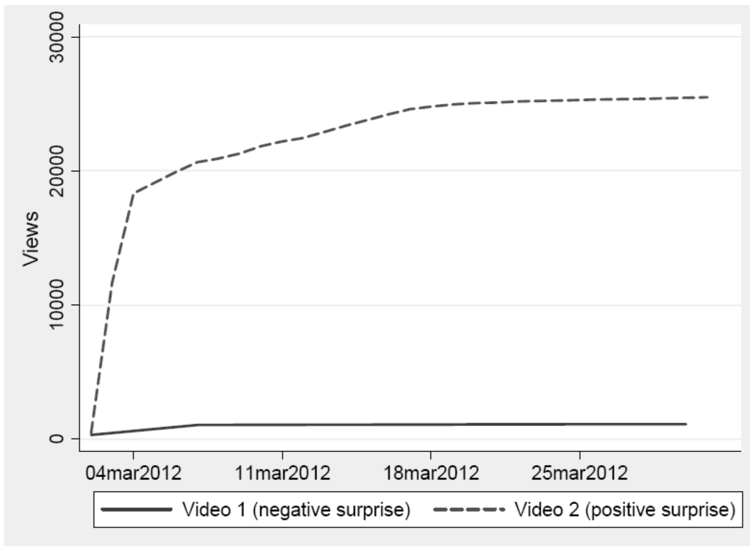

Figure 2.3. Daily Views for Videos with Different Surprises

Figure 2.3 shows a clear example of videos with different surprises. The figure plots the daily video views for a video with a positive surprise (video 2) and a video with a negative surprise (video 1). These two videos belong to the same YouTube video category and have similar initial views, but experience different growth patterns: Video 2, having positive surprise, has a significantly higher growth rate than video 1, having a negative surprise. The first-day views of video 1 and video 2 are roughly the same (304 and 449 respectively). However, at the end of our sample period, views of video 1 and views of video 2 are 1,102 and 25,508 respectively. As was stated in our theoretical model, this striking difference can be caused by social learning over time.

2.5.2 A Test of Social Learning

Hypothesis 2.1 indicates that in the presence of social learning, the growth rate of views of a video with a positive surprise is higher than the growth rate of views of a video with a negative surprise. If only network effects exist, we should see no significant difference between the growth rate of a video with a positive surprise and that of a video with a negative surprise.

Our empirical approach is based on the literature on treatment effects (Wooldridge 2007). The positive surprise is interpreted as the "treatment", and views of "treated" videos are compared to the views of videos with negative surprises.

To test whether the difference between the growth rates is significant, we estimate the following model using difference-in-difference:

$$\ln views_{jt} = b_0 + b_1 t + a_j + \Psi_{jt}' b_2 + b_3 (t \times D_j) + \mu_{jt}, \tag{2.3}$$

where $\ln views_{jt}$ is the log of views of video j at time t, t represents the time period, a_j represents unobserved fixed effect of video j, and Ψ_{jt} includes the characteristics of video j that change over time, such as $rating$ (the video ratings), $favs$ (the number of YouTube Favorites), $comment$ (the number of video comments), $lvideo$ (the log number of uploaded videos of the channel), $lvviews$ (the log of total video views of the channel), and $subs$ (the number of channel subscriptions). We control the marketing efforts of YouTube providers on Twitter, which are measured by sum_upload_{jt}, the total number of tweets containing the unique YouTube video ID and the word "uploaded." Unlike Duan et al. (2009) and Susarla et al. (2011), equation (2.3) does not contain lags of accumulative views, because the lag terms do not help distinguish between social learning and network effects. In our context, there are two reasons why lags of views can have a positive effect on current views: (1) Consumers learn from other people's choices.

They infer the quality is high when they see a larger number of accumulative views. (2) Consumers can obtain a higher utility from a larger view base because they enjoy discussing a video with their peers.

In the regression, D_j is a dummy variable indicating whether the surprise of video j is positive ($D_j = 1$, if the surprise is positive; $D_j = 0$, otherwise), and μ_{jt} is the error term. Following the literature on treatment effects (Wooldridge 2007), we make the unconfoundedness assumption: $t \times D_j$ is strictly exogenous. Note that correlation between $t \times D_j$ and μ_{jr} for any time t and time r causes inconsistency in regression coefficients. Thus, we need to control the time-varying heterogeneity (Ψ_{jt}), and the unobserved fixed effect in the regression. If the surprise assignment (positive or negative) changes in reaction to past outcomes on $\ln views_{jt}$, strict exogeneity can be violated (Wooldridge 2007). However, the surprise assignment is determined at time 1 and is independent of the idiosyncratic views shocks in period t. So strict exogeneity is a reasonable assumption.

We are interested in the difference-in-difference estimator, b_3. If $b_3 > 0$, then the difference between the growth rates is positive. It is consistent with Hypothesis 2.1 and implies the existence of social learning. If $b_3 = 0$, then the growth rates of video views with different surprises are the same, which indicates there is no significant social learning on YouTube.

Regression results are shown in Table 2.4. In the table, $interaction = t \times D_j$. Column 1 shows the results from a regression that includes all the coefficients specified by (2.3). In this regression, $\hat{b}_3 = 0.0113$ and is significantly positive, which confirms Hypothesis 2.1. Column 2 shows that the basic result is robust to an alternative model specification.

Table 2.4: Regression of Video Views on Surprises: A Test of Social Learning

	(1)	(2)	(3)	(4)	(5)
interaction	0.0113***	0.0113***	0.0280***	-0.00118	0.0176***
	[12.58]	[12.55]	[7.727]	[-0.632]	[30.86]
rating	0.262***	0.268***	1.542***	1.105***	0.206***
	[2.934]	[2.997]	[2.771]	[2.652]	[2.442]
favs	0.000110***	0.000125***	0.000834***	0.000134	0.000172***
	[2.920]	[3.810]	[6.157]	[0.890]	[4.87]
comment	0.000213***	0.000222***	-0.000406***	0.000423***	0.000134***
	[7.835]	[8.589]	[-3.369]	[4.506]	[5.204]
sum_upload	0.0300				0.0241
	[0.461]				[0.390]
Observations	9060	9060	3390	1290	9060
R-squared	0.328	0.328	0.459	0.564	0.110

t-statistics in brackets: *** *p<0.01*, ** *p<0.05*, * *p<0.1*

The effect of social learning may differ across video categories. Column 3 and Column 4 are regression results for subsamples of two different categories: "Music" and "Tech." According to a survey by Sysomos Inc.[12], music is the most popular category on YouTube, and Tech is the least popular category. In Column 3, $\hat{b}_3 = 0.0280$, which is significantly larger than the estimate in Column 1 and 2. It shows that social learning is more pronounced for videos belonging to category "Music." In Column 4, $\hat{b}_3 = -0.00118$, which is not significantly different from zero. Social learning is not significant for videos in category "Tech." The evidence shows that social learning is more pronounced for videos belonging to more popular categories. Given the large number of reviews for videos belonging to more popular categories, consumers rely more on social learning in the context of YouTube. Column 5 shows that the results are robust when we define the surprise as the difference between realized video ratings and predicted video ratings. These results help social media managers answer the question

[12] http://www.sysomos.com/reports/youtube/

regarding which content creators should be subsidized. We provide evidence suggesting that social learning is more pronounced for videos in more popular categories. The role of social learning becomes more salient in a mass market than in a niche market. As such, content creators in popular categories may benefit more from subsidizing.

Hypothesis 2.2 indicates that social learning is more important for videos with less precise priors. To test the hypothesis, we estimate the following model:

$$\ln views_{jt} = b_0 + b_1 t + a_j + \Psi_{jt}' b_2 + b_3 (t \times D_j) + b_4 (t \times prior_j)$$

$$+ b_5 (t \times D_j \times prior_j) + \mu_{jt}, \qquad (2.4)$$

where $prior_j$ is a measure of the prior precision of video j. Here we propose the total views of the provider's channel page on the first day (March 2, 2012) to empirically identify which videos have more precise priors. YouTube users upload videos to their YouTube channels. A consumer has a better idea of the quality of a new video published by a high-ranking channel (in terms of channel page views) because the consumer is more likely to have watched another video published by the same channel before. In this case, videos published by higher-ranking channels have more precise priors. We divide the sample into two equally sized groups by channel views: the high-ranking group and the low-ranking group. If a video belongs to the high-ranking group, then the dummy $prior_j = 1$; otherwise, $prior_j = 0$.

The coefficient of interest here is b_5 in the regression model (4). Table 2.5 illustrates that the empirical evidence is consistent with Hypothesis 2.2. In the table, $tprior = t \times prior_j$, and $tdprior = t \times D_j \times prior_j$. In Column 1, we find that the coefficient on $tdprior$, $\hat{b}_5$, is -0.0151, which is significantly negative. We can consider two identical videos with the same positive surprise except for the fact that the first belongs to the high-ranking group and the second belongs to the low-ranking group. In the presence of social learning, a negative b_5 implies that the growth rate of views of the second video is higher than the growth rate of views of the first video. In other words, social learning has a greater effect on videos with less precise priors, which supports Hypothesis 2.2. Column 2 indicates that the estimate of b_5 is robust to a different specification.

Table 2.5: The Effect of Prior Precision on Social Learning

	(1)	(2)	(3)
interaction	0.0152***	0.0333***	0.0129***
	[12.15]	[9.581]	[10.76]
tprior	0.0113***	0.0482***	0.0236***
	[15.15]	[25.47]	[30.88]
tdprior	-0.0151***	-0.0523***	-0.0162***
	[-8.142]	[-10.16]	[-9.050]
rating	0.154*		0.0371
	[1.744]		[0.445]
favs	0.000136***		6.40e-05*
	[3.677]		[1.823]
comment	0.000160***	0.000871***	0.000265***
	[5.943]	[18.99]	[10.43]
sum_upload	0.0157		0.186***
	[0.246]		[3.058]
Observations	9060	9060	9060
R-squared	0.353	0.253	0.222

t-statistics in brackets, *** *p<0.01*, ** *p<0.05*, * *p<0.1*

We also show that the regression results are robust to different measures of the prior precision. We use another measure of the prior precision: the number of subscribers. If more YouTube users subscribe to a channel, consumers are more certain about the quality of the videos from that channel. Similarly, we divide the sample into two equally sized groups, based on subscribers rank. If a video belongs to the high-ranking group, then the dummy $prior_j = 1$; otherwise, $prior_j = 0$. Column 3 in Table 2.5 shows that the coefficient on $tdprior$ is significantly negative. It implies that different measures of the prior precision do not affect our key results.

2.5.3 A Test of Network Effects

In this section, we test the existence of network effects on YouTube using the presence of in-stream ads as a source of exogenous variation for existing levels of video views. YouTube in-stream ads run only on partner videos. Only successful content creators are qualified for the partner program, and videos published by them might contain in-stream ads. It is reasonable to assume that the presence of in-stream ads is exogenous in our context. One might argue that advertisers are more likely to use the channels that have higher viewership and higher quality videos. However, all of our sample videos are published by the top content creators on YouTube, and almost all of them have been partners before our data collection. The presence of in-stream ads is not likely to be correlated with video quality.[13]

Network effects are identified by isolating the surprises caused solely by the presence of in-stream ads. YouTube bundles video content with in-stream ads, which are intrusive to many consumers.[14] The presence of in-stream ads is a negative shock that can reduce the viewer base. If network effects exist on YouTube, the negative shock lowers the growth rate of views at time 1, and then it further lowers the growth rate at time 2. As time goes on, we should see a significantly negative self-reinforcing feedback loop. However, such a negative shock does not contain any information of the video quality. If

[13] Empirically, we find that the presence of in-stream ads is not significantly correlated with viewership in our sample.

[14] Wilbur (2008) estimates a two-sided model of the television industry and shows that viewers tend to be averse to advertising. Anderson and Gans (2011) study an advertising-sponsored content provision model and interpret advertising clutter as a "price" paid by viewers who enjoy the content.

social learning is the sole form of social contagion, the Bayesian learning process remains unchanged. The negative shock can decrease the viewership at time 1, but the long run growth rate of video views should not be affected significantly (no self-reinforcing feedback loop).[15] If a consumer is shown an ad before the video, one may think this could impact social learning in that this would result in lower consumer satisfaction and more negative word of mouth. However, the additional information about ads from the peers is redundant. When consumers make decisions on whether to watch the video, they know whether the video contains an ad.[16] In summary, if there exists only social learning with network effects absent, the presence of ads is a transitory shock that does not have significant long run effect. If network effects exist, the presence of ads results in a negative self-reinforcing feedback loop.

We re-estimate (2.3) to test network effects, using two-stage least squares (2SLS) regression. We instrument the surprise dummy D_j using the in-stream ads ads_j. ads_j is a dummy, where $ads_j = 1$ if the video has an in-stream ad, and $ads_j = 0$ otherwise. Generally, 2SLS is used to avoid endogeneity. However, we use the first-stage regression to isolate the surprises that are caused solely by the shock of in-stream ads. We are interested in the coefficient b_3 in the regression model (3) in Section 2.5. If Hypothesis 2.3 is supported, we expect to see that $b_3 > 0$, which implies that negative surprises lower the future views. In this case, network effects are significant. If $b_3 = 0$, it means that the network effects on YouTube are not significant, and the dominant causal social contagion is social learning.

---

[15] Let video *j* is a video without an ad, and video *j'* is a video contains an ad. Other things being equal, we can obtain $\ln views_{jT} > ln\, views_{j\prime T}$, $\ln views_{jT-1} > \ln views_{j\prime T-1}$, and $\ln views_{jT} - \ln views_{jT-1} \approx \ln views_{j\prime T} - \ln views_{j\prime T-1}$ from our model of social learning.

[16] When consumers are shown an ad before the video, they can choose to switch to other videos. It is equivalent to not watching the video.

Table 2.6 shows the results of the test. b_3 is the coefficient on $interaction$ ($t \times D_j$). Column 1 and 2 represent different model specifications. We find that $\hat{b}_3$ is significantly positive under all specifications. The result suggests that social learning is not the only causal social contagion on YouTube, and network effects also play a critical role. The test confirms Hypothesis 2.3: the existence of network effects on YouTube. Column 3 shows that the results are robust when the surprise is defined as the difference between realized video ratings and predicted video ratings.

Table 2.6: A Test of Network Effects: 2SLS

	(1)	(2)	(3)
interaction	0.00670**	0.00667**	0.0036***
	[2.540]	[2.509]	[2.583]
rating	0.436***	0.438***	0.430***
	[5.213]	[5.248]	[5.210]
favs	0.000170***	0.000174***	0.0002***
	[4.794]	[5.587]	[5.582]
comment	0.000138***	0.000141***	0.0001***
	[5.425]	[5.853]	[4.831]
sum_upload	-0.00957		-0.0198
	[-0.157]		[-0.334]
Observations	9060	9060	9060
R-squared	0.241	0.241	0.241

t-statistics in brackets, **** p<0.01, ** p<0.05, * p<0.1*

A straightforward implication of our study is that YouTube should take social learning and network effects into account when fostering the growth of video views. Our results suggest that social contagion on YouTube is driven by both social learning and by network effects. Considering that the amount of a traditional marketing campaign of YouTube content is limited, consumers rely heavily on advice from others to make decisions about watching videos. Social learning and network effects differ from traditional marketing activity in their social multiplier effect. From a managerial perspective, YouTube can play a much greater role in encouraging the creation of original content because, as it nurtures and subsidizes individual content creators, the multiplier effect of both social learning and network effects extends their reach.[17]

2.6 Application: How to Go Viral?

Social learning and network effects outline two ways that a video could go viral and gain success. Through examining the most popular videos on YouTube, we are able to categorize them into two distinct groups: the group consists of videos that feature high quality, engaging scenes, articulated story lines (high-quality videos), and the other group of videos often include questionable behaviors that deviate from social norms yet still gain tremendous popularity (attention grabbers). The recent "Pussy Riot" incident in Russia serves as a good example of a typical attention grabber. This Russia-based feminist rock band protested against the political scene in Russia through unorthodox musical performances and produced YouTube videos that went viral overnight. It is worth noting that Pussy Riot did not gain international fame through their musicality per se; instead, most viewers were drawn to those videos out of curiosity and were interested in the messages the music carried.

One type of strategy often adopted is the inclusion of controversial elements in videos. Instances of such often provoke controversy and stir heated discussion revolving around those contents. This type of videos tends to attract critical reviews from both sides of the spectrum; viewers feel strongly and emotionally attached to the video in either extremely positive or negative ways. In contrast to those quality-oriented productions, the

[17] In fact, YouTube is providing creators with resources and opportunities to improve their skills, build larger audiences, and make more money through its partnership program. As a New York Times article reported (Miller 2011), a sketch comedy show called "AsKassem," is financed by grants from YouTube. The amount of content on YouTube covered by partnership agreements has risen steadily, to 10% of the total videos.

goal of attention grabbers is primarily to attract attentions or promote ideas. Intuitively speaking, we would not expect too much social learning effect to take place for the popularity of this type of video. In an analytical model, Eliaz and Spiegler (2011) shows that a firm can earn higher profits by employing pure attention grabbers with positive probability. Similarly, we propose that, as suggested by their discussion-provoking nature, videos with attention grabbing content can initiate higher network effects, and viewers find it valuable to watch them because these videos allow them to engage in discussions with their social contacts. Therefore, we hypothesize that this type of videos gains popularity mostly through network effects as opposed to social learning:

Hypothesis 2.4. *(a) Network effects are more pronounced for videos with attention grabbing content. (b) Social learning is more pronounced for high-quality videos.*

To test Hypothesis 2.4, we estimate the following two models:

$$\ln views_{jt} = b_0 + b_1 t + a_j + \Psi'_{jt} b_2 + b_3 (t \times D_j)$$

$$+b_4 \left(t \times attention_j\right) + b_5 \left(t \times D_j \times attention_j\right) + \mu_{jt}, \qquad (2.5)$$

and

$$\ln views_{jt} = b'_0 + b'_1 t + a_j + \Psi'_{jt} b_2{}' + b'_3 (t \times D_j)$$

$$+b'_4 \left(t \times quality_j\right) + b'_5 \left(t \times D_j \times quality_j\right) + \mu_{jt}, \qquad (2.6)$$

where $attention_j$ is a measure indicating whether video j is a video with attention grabbing content, and $quality_j$ is a measure indicating whether video j is a high-quality video. Here we use views rank and rating rank to empirically identify videos with high quality or attention grabbers. We define high-quality videos as videos with both high views rank and high ratings, and attention grabbing videos as videos with high views rank but mixed ratings. The co-existence of extremely high and extremely low ratings often suggests controversy. Specifically, if both the views and the rating of video j at the end of our sample period rank among top 25% of total videos, then it is considered as a high-quality video, and the dummy $quality_j = 1$; otherwise, $quality_j = 0$. If the views of video j at the end of our sample period rank among top 25%, but the rating is in the lowest 25%, then it is a video with controversial content, and the dummy $attention_j = 1$; otherwise, $attention_j = 0$.

Table 2.7. High-Quality Videos vs. Attention Grabbers

	(1)	(2)	(3)
t * D * attention	0.0145***		
	[6.330]		
t * D * quality		0.368**	-0.0208
		[2.014]	[-0.281]
rating	0.0307	0.0525	0.188
	[0.393]	[0.238]	[1.295]
favs	0.000145***	-0.000621*	0.00034***
	[4.416]	[-1.672]	[4.309]
comment	0.000140***	0.000853***	1.50e-06
	[5.821]	[10.21]	[0.0169]
sum_upload	0.0342	0.0715	-0.164
	[0.603]	[0.378]	[-1.254]
Observations	9060	9060	9060
R-squared	0.193	0.117	0.161

t-statistics in brackets, *** *p<0.01,* ** *p<0.05,* * *p<0.1*

In order to study the network effects for different videos, we estimate the regression model (2.5) using 2SLS. Similarly, we instrument the surprise dummy D_j using the in-stream ads ads_j. b_5 in (2.5) is the coefficient of interest. $b_5 > 0$ means the impact of network effects is larger for attention grabbing videos. Column 1 of Table 2.7

presents the regression results. We find that b_5 is significantly positive, supporting Hypothesis 2.4(a).

To study social learning, we first estimate regression model (2.6) without using an instrument variable. In Column 2 of Table 2.7, b'_5 is the coefficient on $t \times D_j \times quality_j$. We find that b'_5 is significantly positive, which suggests that social contagion is more pronounced for high-quality videos. However, it does not provide sufficient evidence for social learning because social contagion can be driven by network effects as well. Therefore, we re-estimate (2.6) using ads_j as an instrument variable. The result is presented in Column 3 of Table 2.7. We find that $b'_5 = -0.0208$ and is not significant, which indicates that high-quality videos do not have higher network effects. Combining the results shown in Column 2 and 3, we can conclude that social learning is more pronounced for high-quality videos. In other words, Hypothesis 2.4(b) is also supported. Our empirical results of hypothesis testing provide supports for the strategic use of attention grabbers (Eliaz and Spiegler 2011).

2.7 Conclusions

Previous literature has focused on distinguishing social contagion from homophily, but only provides limited insights into how to disentangle social learning and network effects in the context of UGC. In this chapter, we developed an empirical framework that allows us to make a causal inference about the presence of social learning and network effects on YouTube. More specifically, by applying a theoretical model, we conducted a two-stage test and examine the existence of social learning and network effects using a unique data set from YouTube.

Although in this study we only focused on social learning and network effects on UGC sites, our tests are relatively generalizable and can be practically carried out by practitioners in social media. We categorized the most popular videos on YouTube into quality-oriented videos and anti-social videos, and found that videos with anti-social content initiate higher network effects than quality-oriented productions. These findings provide a nuanced view of how YouTube can best subsidize the content creators.

While this study has highlighted the importance of social learning and network effects, our work does not consider the effect of network characteristics and network topological structure on social contagion (Ghose et al. 2012). Further work could incorporate network data to examine the effect of network structure and tie strength on social learning and network effects.

Chapter 3: Improving Internet Security Through Social Information and Social Comparison

3.1 Introduction

2011 was a busy year for cyber attacks on many organizations, with targeted attacks increasing by 400%. Industries such as credit card companies, gaming platforms, banks, retailers, TV networks, and government agencies all fell victim to cybercrime, which is not only increasing in frequency but also in the severity of damage. According to the Ponemon Institute, the median cost caused by cybercrime is $5.9 million per year per company, with a range from $1.5 million to $36.5 million. The costs consist of both direct expenses (recovery, detection, etc.) and indirect costs (information loss, business disruption, revenue loss, equipment damages, etc.). However, the study by Ponemon Institute also shows that nearly all of these attacks were avoidable. Most attacks were carried out using fairly simple methods and could have been stopped easily with basic or intermediate controls. Although most attacks were targeted, the target selection was based more on opportunity than on choice. Most organizations fell victims not because they were pre-identified but because they were found to possess exploitable vulnerabilities. About 50-75% of security incidents originated from within an organization (D'Arcy et al. 2009). Ninety-six percent of victim organizations subject to the Payment Card Industry Data Security Standard (PCI DSS) were not in compliance.

Organizations generally underinvest in Internet security because of the following reasons. First of all, Internet security is often considered too expensive to achieve. Security products and services are sometimes regarded as useful and desirable, yet not affordable. High-level security practices can be reinforced to prevent security disasters and control the damage. The deployment of such practices, however, is a costly endeavor for organizations without assured significant benefit. With the proliferation of mobile devices, the increasing number of locations and devices where information can be stored further adds to the cost for prevention and protection. Second, although the costs for security are too high, the rewards are unclear. It is difficult to measure the risk and potential costs of security breaches beforehand. The frustrating fact about security is that although insecurity is easy to prove, security can never be conclusive. Third, the absence of legislative enforcement leads to the lack of transparency. Although recent progress in data breach notification laws requires companies to notify those customers whose information has been lost or stolen, companies generally can choose not to reveal publicly any attacks, in order to avoid reputation loss. Without transparency, organizations can claim to be secure even if their systems are, in fact, vulnerable, and customers cannot accurately estimate the risk of doing business with them. Moreover, Internet security is a public good in that an organization's security (insecurity) can benefit (hurt) others. The security vulnerabilities of an organization are often used against other organizations. For example, botnets opportunistically scan the Internet to find and compromise systems with exploitable weaknesses. These compromised computers are then utilized to collectively attack other targeted systems as in a typical denial of service attack.

Although they focus on technical solutions, existing studies often tend to ignore the motivational issue, which is a common problem in private provisions of public goods such as charitable giving (Frey and Meier 2004, Shang and Croson 2009) and contribution to online communities (Bulter 2001, Chen et al. 2010). Social psychologists have documented the existence of social loafing—that people exert less effort on a collective task than they do on a comparable individual task. According to social comparison theory, which was initially proposed by Festinger (1954), people have the desire to gain information on others and evaluation on themselves. When information on others is available, people tend to evaluate themselves in comparison with others. As a result of the self-evaluation, the existence of discrepancy in a social group would lead to action on the part of group members to reduce the discrepancy. People generally care about their social status, often measured by ordinal ranks within their social groups, especially when status is made public and can influence one's reputation (Griskevicius et al. 2010). As a result, status competition is often utilized to encourage desirable behaviors. Because of the concern for customer switch, organizations have even stronger incentives than individuals to maintain their status among peers. Reputation in Internet security signals a company's valuation for customer information and ability to take appropriate security controls. In the present article,, we propose to solve the underinvestment problem by making such information publicly available, in order to solicit social comparison and status competition. Equivalent to rewarding prosocial behavior with status and prestige, we can penalize proself behavior with shame and reputation loss by making these behaviors notorious.

We incentivize organizations to increase security spending through our reputation system, an online website that regularly aggregates individual organizations' security information and releases explicit comparison results as relative performance ranking to the public. It is worth noting that often it is the information aggregation and feedback rather than the information itself that is missing. In the present study, we make use of the available information through third-party monitoring on outgoing spam as the focus security issue. However, the methodology also applies to other security problems, for which data can be collected through public policies on mandatory reporting, in the absence of available data. It has been recognized that a key factor required to improve Internet security is the gathering, analysis, and sharing of information related to security issues (Gal-Or and Ghose 2005). The Securities and Exchange Commission (SEC) formally asked public companies to disclose cyber attacks against them in October 2011. However, no pre-attack information is currently available for businesses and individuals to take precautionary actions. To solicit social comparison, the social information provided needs to reveal what constitutes the right behavior and who behaves that way and who does not.

To test the impact of the specific information released through our website, we conducted a field quasi-experiment in which the released information was used as experimental treatment. To draw attention to our system, we deliberately chose the United States, Canada, Belgium, and Turkey as four treatment countries and did extensive promotion for our website within these countries. The treatment countries were then matched with four control countries according to population and the severity of the security problem before our experiment. Countries were used as clusters of organizations so that an organization was compared to other organizations within the same country. For organizations in the treatment group, the information on the organizations with the severest security problem in the country was released monthly on our website, whereas similar information was kept internally only for the control group. Although the treatment assignment is at the cluster level, the measurement is at the individual level. The field setting ensures that organizations and the public behave in a natural manner. A difference-in-difference model is used to test the treatment effect. The results show that the treated organizations improved their security situations more than the control organizations. We also find that the more security observed for other organizations, the more likely an organization will be to improve its own security situation.

Our approach for improving Internet security is complementary to existing technical approaches. The vast technical literature, especially in the computer science area, has focused on the development of technologies to secure computer systems, such as secure networking protocols, intrusion detection techniques, database security methods, and access control technologies (Ransbotham and Mitra 2009). By focusing on organizations' incentives to invest in these technologies, we aim to extend prior work and provide a more comprehensive lens for studying Internet security. Our study sheds light on public policy issues concerning security information disclosure and provides a new perspective for dealing with other environmental issues such as pollution, energy conservation, and global warming, where externality leads to a lack of incentives for taking pro-social behavior.

3.2 Literature Review

3.2.1 Internet Security

Existing literature on information security focuses on organizational strategies that can be used for reducing system risk, including deterrence, prevention, detection, and recovery (Straub and Welke 1998). For deterrence and prevention, most previous studies, from the organizational perspective, have examined the impact of security policy and practice on information systems abuse or misuse (Kankanhalli et al. 2003, D'Acy et al. 2009). For detection and recovery, research has been focused on how to identify attack traffic that could originate from both internal and external sources in a cost-effective way (Yue and Cakanyildirim 2007, Mookerjee et al. 2011) Specifically for anti-spam, the filtering techniques consist of machine learning (Goodman et al. 2007), crowdsourcing and IP blacklisting (Ramachandran et al. 2011), screening humans from bots for botnets (Isacenkova and Balzarotti 2011), and Domain Keys Identified Mail (DKIM) (Moyer

2011). However, the problem for any technical solutions is that hackers can always respond strategically. The interplay is an endless cat-and-mouse game.

Security vulnerability disclosure is an area of public policy that has been subject to considerable debate. Studies on software vulnerability disclosure have shown that although disclosing vulnerability information provides an impetus to the vendor to release patches early, instant disclosure leaves users defenseless against attackers who can exploit the disclosed vulnerability. Arora et al. (2004a) found that although vendors are quick to respond to instant disclosure, vulnerability disclosure also increases the frequency of attacks. Arora et al. (2004b) suggested that the optimal vulnerability disclosure depends on underlying factors such as how quickly vendors respond to disclosure by releasing patches and how likely attackers are to find and exploit undisclosed or unpatched vulnerabilities. Although there has been no public disclosure on information security vulnerability, industry-based Information Sharing and Analysis Centers (ISACs), where security breach information is revealed to information-sharing alliance, has been established to facilitate the sharing of security information to enhance and protect critical cyber infrastructure. Gal-Or and Ghose (2005) studied the economic incentives for security information sharing and found that information sharing yields greater benefits in more competitive industries. Gordon et al. (2003) examined how information sharing affects the overall level of information security when firms face the trade-off between improved information security and the potential for free riding.

3.2.2 Regulations on Information Disclosure

Security information disclosure laws have been focused on data breach notification. Although a national data-breach law is still under consideration, as of August 20, 2012, 46 U.S. states and the District of Columbia, Guam, Puerto Rico and the

Virgin Islands have enacted legislation requiring notification of security breaches involving personal information. The specific requirements of notification laws vary across states. Some laws require notification when the personal information is reasonably assumed to have been acquired by an unauthorized party, whereas others require notification only if it is reasonable to believe the information will cause harm to consumers. The consequences of not complying differ from state to state as well. However, the rationales for these laws are consistent, which is also consistent with our rational for public disclosure of security vulnerabilities, that notification can provide public information and create an incentive for all firms (even those that have not been breached) (Schwartz and Janger 2007, Romanosky et al. 2011).

However, the impact of data breach disclosure is still in debate. The concerns include the following: (1) Firms must comply with multiple, disparate, and perhaps conflicting state laws (Romanosky et al 2011); and (2) notifications simply shift the burden to consumers if breaches really cause harm (Cate 2009). Otherwise, they are just unnecessary costs. Romanosky et al. (2011) found that data breach disclosure can reduce identity theft caused by data breaches. Campbell et al. (2003) found a highly significant negative impact of security breaches reported in newspapers on the stock price of the breached company only when the breach involved unauthorized access to confidential data. In contrast, Kannan et al. (2007) found that security breach announcements have no significant negative impact on market return in the long run.

The impact of information disclosure has also been widely studied in areas other than security. Jin and Leslie (2003) studied health information disclosure in the restaurant industry and found that disclosing hygiene quality information increases health inspection scores and lowers certain diseases. Cain et al. (2005) examined the effect of disclosing conflicts of interest and found that the disclosure can have perverse effects because advice receivers do not discount advice sufficiently, and that advice givers exaggerate advice even further. Other information disclosure studies are related to auto safety, public education, and so on (Fung et al. 2007). These studies provide some evidence of how information disclosure can affect firm behavior. On the basis of these studies, we further add the aggregation and presentation of information, which can leverage reputation and peer influence to enhance disclosure effect.

3.2.3 The Economics of Internet Security

It has long been recognized that Internet security is not a problem that technology alone can solve (Arora et al. 2004a). Many security questions are at least as much economic as technical. Fundamentally, Internet insecurity is the result of perverse incentives, which are distorted by network externalities, asymmetric information, moral hazard, adverse selection, liability dumping, and the so-called tragedy of the commons (Anderson 2001). Systems fail often because of misplaced economic incentives: The people who could protect a system are not the ones who suffer the costs of failure (Schneier 2002). Security failure is caused as much by bad incentives as by bad design (Anderson and Moore 2006). Meanwhile, hacking has evolved over the past a few years to become a well-organized, sophisticated underground market.

The economic incentive problem is caused by negative externality of insecurity. Externality happens because social costs or benefits are not equal to private costs or benefits (Pigou 1920, Dahlman 1979). Negative externality happens when social costs are greater than private costs, whereas positive externality happens when social benefits are greater than private benefits. Security vulnerabilities of a system are often exploited by hackers to attack other systems. For example, spam has such an extreme negative externality that the social costs are about 100 times the private benefits (Rao and Reiley

2012). More and more studies have recognized the importance of security externalities and have come up with several economic and legal policy proposals. The standard economic treatment for negative externality is to impose a Pigouvian tax on the activity that generates negative externality (Pigou 1920, Dahlman 1979). For spam, researchers in many studies have proposed to have the spam sender pay the receiver for attention or levy penalties on consumers who purchase goods from spammers (Kraut et al. 2002, Loder et al. 2004). However, these proposals raise the concerns for privacy and account hijacking by hackers. The legal treatment is to let government make law or regulation enforcements. For spam, the legal interventions include requiring legal advertisers to offer opt-in or opt-out choices for email receivers and putting legal pressure on banks that process payments from foreign banks known to act on behalf of spam merchants (Sipior et al. 2004). However, since most security problems such as spam and phishing may involve parties in different administrative areas, jurisdictional boundaries render the proposals unrealistic.

3.2.4 Social Comparison

In a social community, participants tend to compare themselves to others when social information on other participants' behaviors is available, and such social comparisons in turn affect behaviors (Festinger 1954). Perceptions of relative standing can influence many outcomes. A number of studies have found that self-reported happiness may be more sensitive to relative than to absolute income (Hopkins and Kornienko 2004, Luttmer 2005). The interdependent preferences can be represented either by utility functions that depend not only on the absolute value of consumption but also on the average level of consumption within a population, or by including concern for status, the ordinal rank in the distribution. The reasons for status concern may be intrinsic, a fundamental human characteristic, or instrumental: Status is desirable because it allows better consumption opportunities (Hopkins and Kornienko 2004).

The availability of social information is the prerequisite for social comparison. Recent theories on pro-social behavior have focused on "conditional cooperation": People are more willing to contribute when information is provided that many others contribute (Frey and Meier 2004). Satio (2011) suggested that individuals feel ashamed about a choice that does not maximize the payoffs of others only when the choice is made in public. Dillenberger and Sadowski (2010) also proposed that a person's behavior may depend on whether it is observed by someone who is directly affected by it and considered shame as a moral cost for a person's utility. These concepts can be extended to organizational behavior since organizations are concerned about their social image and reputation (Frei 2010), their relative standing in comparison to other businesses. These social factors such as reputation and social image are valuable assets for a business not only because organizations have the desire for prestige, esteem, popularity, or acceptance (Bernheim 1994), but also because they lead to better business opportunities. With the increasing concern for privacy and confidentiality, customers are likely to choose or switch to firms with a more secure information system.

Social comparison and social information are often used to solve the problem of social loafing, the reduction in motivation and effort when individuals work collectively as compared with when they work individually. The reasons include reduced individual motivation and coordination loss (Karau and Williams 1993). Both reasons exist in the context of Internet security. The former is due to the externalities, whereas the latter is due to the cost of security efforts. Reputation loss imposed by making relevant social information available can serve as a binding force against social loafing (Akerlof 1980). Social norm formed through social information provision has two effects on pro-social behavior: the focusing influence whereby norms impact behavior only when an individual’s attention is drawn to them, and the informational influence whereby norms exerts a stronger impact on an individual the more he observes others behaving consistently with that norm (Krupka and Weber 2009). In the present article, we aim to leverage both effects to motivate pro-social behavior.

3.3 Field Quasi-Experiment

Field experimentation has been used extensively to provide solid knowledge of causation for policy evaluation (Duflo et al. 2010). It has also been used to study information security and privacy (Hui et al. 2007). Experimental studies randomly assign participants into treatment groups or control groups. Randomization, although more desirable in an ideal environment, is inappropriate given our circumstance. In the present study, we aimed to evaluate whether public information disclosure can lead to security improvement; thus, the attention to the disclosed information is critical. Rather than randomly choosing some countries for treatment, it is more pragmatic to focus on on the countries where the new information is more likely to receive attention. As a result, we used a quasi-experiment with intentional treatment on North American and European countries, to resemble the randomized field experiment, considering the authors’ PR connections and promotional activities for our website. Quasi-experiments typically occur in real-world settings that more closely resemble the actual contexts and constraints faced by policymakers and practitioners. Although randomized experiments generally have better internal validity (evidence of causation), quasi-experiments often turn out to have better external validity (generalizability).

3.3.1 Outgoing Spam

Internet security is a very broad and general concept that has many dimensions. In the present article, we look into outgoing spam as one specific security issue. Referred to as unsolicited bulk emails, most spam messages are sent by botnets, a collection of compromised computers (bots), without the awareness of the legitimate computer owners. Anti-spam blocklists have spam traps scattered across the Internet and can recognize similar messages received at multiple locations. An estimated 88% of daily worldwide email traffic is spam. Inbound spam refers to the spam received, and many organizations are well equipped to filter spam out of incoming emails before these emails reach their employees or users. However, they have very limited techniques to prevent outbound spam originated from computers within the organizations. Outgoing spam is typically generated via zombie computers, compromised user accounts, or spammers who knowingly abuse their accounts (e.g., in snowshoe spam), and it is a common symptom of more damaging security problems. The same vulnerabilities that enable spam are also openings for other exploits. For example, miscreants can steal existing accounts by tricking end-users (through phishing or by human engineering) into providing their email usernames and passwords. Such stolen accounts can then be used to install botnet spamming malware or other exploits such as Distributed Denial of Service (DDoS) software or sniffers, causing theft of customer records and intellectual property, fraudulent use of corporate online banking, or even employee blackmail.

It is costly to deal with outbound spam, which often leads to major side effects such as IP blocking by RBL, DNSBL, and IP reputation systems. These side effects cause queue buildup on the affected mail server, delays in message delivery, and may result in lost messages and calls from unhappy end-users. They also lead to compromised user accounts and blocking of legitimate outbound email, which then cause damage to reputation, customer relationship, business operation, and eventually lower profit. Unfortunately, conventional anti-spam measures may not work well for outbound traffic. Spam has an extreme negative externality in the sense that the ratio of external costs to private benefits is as high as 100:1, as compared with about 0.1 for pollution and 7:30 for nonviolent property crime (Rao and Reiley 2012).

Therefore, if ISPs are constantly sending out spam, they not only risk being attacked themselves, but also increase the risk faced by other Internet users. In other words, the efforts of reducing outgoing spam can produce a remarkably large positive externality on other users. For instance, in 2011, Microsoft, Pfizer, FireEye network security, and security experts at the University of Washington collaborated to take down Rustock, the largest botnet on record. The takedown of this single botnet was followed by an immediate one-third reduction in global email spam (Rao and Reiley 2012). Hence, outgoing spam is a typical Internet security problem that lacks transparency, costs a lot to deal with, and generates negative externalities. If our approach proves effective in reducing outgoing spam, it can also be used for improving other security dimensions.

3.3.2 SpamRankings.net

We launched a website named SpamRankings.net in May 2011 and have since used it to release country-specific outbound spam information. This website serves as our main instrument to study public security information disclosure and presentation. It displays monthly outbound spam volume and rankings for sample organizations in the treated countries, including the United States, Canada, Belgium, and Turkey. To generate such information as treatment, we gathered and processed a large amount of daily spam data from two blocklists, the Composite Blocking List (CBL), and the Passive Spam Block List (PSBL). The CBL gathers its source data from very large mail server installations and lists IPs exhibiting characteristics that are specific to open proxies of various sorts and dedicated Spam BOTs that have been abused to send spam, worms/viruses. The PSBL is an easy-on, easy-off blacklist that does not rely on testing and has a lower probability of false positives because any user can remove their ISP's mail server from the list.

The raw data include observed spamming IP addresses, corresponding outbound spam volume, and botnet tags in the forms of text files from CBL and Network News Transfer Protocol (NNTP) messages from PSBL. One important step in data processing is mapping IP addresses to netblocks and, subsequently, Autonomous Systems (ASes), which are groups of IP addresses owned by an organization. Organizations with very large networks may use multiple ASes as a way to divide their networks. Therefore, ASes, even within the same organization, are relatively independent of each other. Therefore we use ASes as the measurement level. An AS can be identified by a unique Autonomous System Number (ASN). Lastly, we aggregate the daily outbound spam data into monthly data and derive rankings for each country. We receive more than eight million records per day from CBL and PSBL, which we summarize into about two million spam messages observed from worldwide IP addresses. On the ASN level, we have seen 27,500 ASNs with spam volume over the lifespan of this project. The ASNs are then grouped and ranked by country. The Top 10 organizations with the most spam are listed on SpamRankings.net (Figure 3.1). For each Top 10 ASN, we display the following information: rank, rank in the previous month if it was listed in the previous month ("-" if not), name and website of the organization, ASN, and outgoing spam volume.

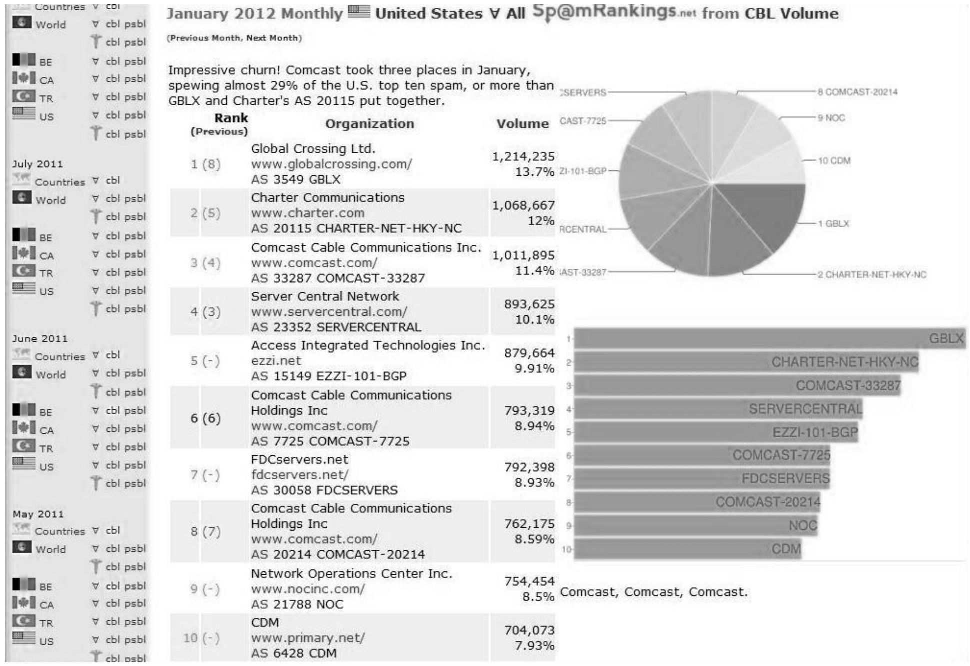

Figure 3.1: Screenshot of SpamRankings.net

3.3.3 Quasi-Experimental Design

To evaluate the impact of spam information released on SpamRankings.net, we used a between-subjects quasi-experimental design with two conditions: the treatment with information released on SpamRankings.net every month, and the control with information kept internally. To strengthen social comparison and reputation concern, information released in the treatment condition is relative ranking with respect to the outgoing spam, which also means that the intervention is at the cluster, the country level. Individual ASN level assignment would have resulted in less meaningful ranking information and weaker social comparison. Therefore, we nested ASNs within countries and assigned countries as clusters of ASNs to the two conditions. Considering the publicity of SpamRankings.net, we specifically chose the United States, Canada, Belgium, and Turkey as four treatment countries. We then matched the treatment countries with four control countries accordingly, based on population and total outgoing spam volume before our experiment, as shown in Table 3.1. Therefore, the control group consists of ASNs in Indonesia, Malaysia, Netherlands, and Iran.

Table 3.1: Country Pairs

Pair	Country	Population	Spam*	Group
1	United States (US)	310,232,863	57,176,031	treated
	Indonesia (ID)	242,968,342	94,435,116	control
2	Canada (CA)	33,679,000	4,387,388	treated
	Malaysia (MY)	28,274,729	6,695,830	control
3	Belgium (BE)	10,403,000	3,781,796	treated
	Netherlands (NL)	16,645,000	6,283,101	control
4	Turkey (TR)	77,804,122	14,759,146	treated
	Iran (IR)	76,923,300	13,291,908	control

* All the data on spam in this chapter refer to the data on outbound

For the treatment group, we made the monthly spam rankings available through SpamRankings.net from May 2011 to January 2012. We treated different countries with differing starting points in time, with the United States in May 2011, Canada in June 2011, and Belgium and Turkey in July 2011. This sequential release was designed to accumulate publicity for our ranking site before getting into the full-scale experiment. For the control group, we did not publish any information on outbound spam, but the same data were collected and kept internally. We also collected static information on each AS, including number of IP addresses, number of unique IP addresses, number of prefixes, number of regions, network name, website, network type, traffic level, inbound versus outbound traffic ratio, and geographic scope. The primary outcome of interest is the outgoing spam volume. The sample ASes were included in either the treatment or control condition because they were observed to send out spam during May 2011 to

January 2012. Therefore, we have a selection bias toward ASes with more severe outgoing spam problems. In the specific context of the present study, this was not a problem since these ASes were the ones we intend to impact. All of the ASes remained through the entire experiment.

Since we wanted to engage both organizations and consumers and observe their natural reactions, it was critical for the success of the experiment to accumulate sufficient visibility and attention of SpamRankings.net. We promoted the website through different channels, including social media such as YouTube, Twitter, and blogs, conferences, and press releases, to increase its impact. We also received much feedback and collaboration requests from industries and observed that some organizations had already tried to take their names off the list on SpamRankings.net. For example, we received the following comment from a Chief Security Officer of a medical center, which also confirms that some outgoing spam could be reduced using basic controls:

> The first time we were rated #1 on your list, we noticed that one of our users had generated thousands of spam messages and asked her to change her password—that stopped the spam immediately. The next month, we found another user who had just given up her credentials and got her to change her password as well. I spoke with a colleague at one of the other medical centers ranked on your site and he mentioned they have the same problem…The listing on your site added additional impetus to make sure we 'stay clean' so in that regard, you are successful.

3.4 Data

We collected outbound spam data on the top 250 most spamming ASNs each month from March 2011 to January 2012 for the eight selected countries. Table 2 shows the summary statistics of observed sample ASNs by country. Only the United States had over 250 spammers for some months, but the top 250 ASNs accounted for over 95% of the total outbound spam. The total unique sample size was 1,718 ASNs, with 1,177 ASNs in the treated group and 541 ASNs in the control group. However, if we look at the average number of ASNs with observed outgoing spam per month, we have a more balanced treatment group and control group. The unbalance is the result of the observation that spamming ASNs for the treatment group varied significantly from month to month, indicating that ASNs in the treatment group reduced their outbound spam more quickly than those in the control group. Results for average maximum of spam percentage by ASN show that a few ASNs were responsible for the most outbound spam volume in Indonesia, Turkey, Belgium, and Malaysia. Especially for Indonesia, the most spamming ASN sent out 83.46% of total spam on average. However, for the United States, the most spamming ASN accounted for only 6.89% of total spam on average.

Table 3.2: Observed Sample ASNs by Country

	Number of ASNs	Average number of ASNs with positive spam volume per month	Average max of spam volume by ASN	Average max of spam percentage by ASN
Treated				
US	699	250	3,414,080	6.89%
CA	316	175	1,261,576	20.94%
BE	56	31	1,116,462	44.70%
TR	106	63	5,051,160	48.08%
Sum	1177	519		
Control				
ID	229	190	45,903,492	83.46%
MY	57	44	1,958,958	43.11%
NL	170	101	1,067,763	22.91%
IR	85	77	2,575,711	27.89%
Sum	541	413		

Table 3.3 Outgoing Spam Volume Observed Per Month Per Country

	Pretest Period*	Test Period*	Difference	Percentage
Treated group				
US	105,347,424	33,007,389	72,340,035	68.67%
CA	7,786,736	3,949,362	3,837,374	49.28%
BE	3,812,537	1,663,925	2,148,612	56.36%
TR	14,758,174	8,052,961	6,705,213	45.43%
Average	32,926,218	11,668,409	21,257,809	54.93%
Control group				
ID	93,416,115	46,320,078	47,096,037	50.42%
MY	6,361,998	3,684,334	2,677,663	42.09%
NL	7,261,624	2,086,952	5,174,672	71.26%
IR	10,590,092	8,515,070	2,075,021	19.59%
Average	29,407,457	15,151,609	14,255,848	45.84%

*****For US, the pretest period is 3/2011-4/2011; the test period is 5/2011-1/2012.
For CA, the pretest period is 3/2011-5/2011; the test period is 6/2011-1/2012.
For other countries, the pretest period is 03/2011-06/2011; the test period is 7/2011-1/2012.

Table 3.3 summarizes the outbound spam volume by country for both the periods before (Pretest Period) and during (Test Period) the experiment. This comparison presents the average outbound spam volume per month and the difference. On average, the outbound spam volume of the four countries in the treated group dropped by 54.93%, whereas the number for the four countries in the control group was 45.84%.

3.5 Statistical Models

We estimate a linear model to test the effect of security information disclosure. First, we employ a simple difference specification to directly compare the treatment and control groups:

$$Y_{ict} = \theta_0 + \theta_1 D_c + \varepsilon_{ict} \ , \tag{3.1}$$

where the dependent variable Y_{ict} is the outcome of interest for AS i in country c at time t during the experimental period, and D_c is an treatment indicator variable for whether country c received security information disclosure, which remains invariant during experiment period for country c. Hence, the estimate of the coefficient θ_1 indicates the difference between treatment and control countries. We utilize this model to compare baseline differences in pre-treatment conditions and to test the effect of spam information disclosure on firms' outbound spam.

Since the assignment of countries to treatment and control groups is not random in the present study, the outbound spam is likely to be affected by pre-treatment conditions. It is thus necessary to include observable AS characteristics and baseline spam volumes as control variables in (3.1) to improve the precision of the estimated treatment effect. Therefore, we also estimate the following specification:

$$\mathrm{Y}_{ict} = \theta_0 + \theta_1 \mathrm{D}_c + \theta_2 \mathrm{X}_{ic} + \omega_p + \varepsilon_{ict}\,, \tag{3.2}$$

where Y_{ict} and D_c are defined as in (3.1), and X_{ic} is a vector of pre-treatment AS characteristics including baseline spam volume and number of IP addresses. Since the assignment to treatment and control groups was stratified within country pairs (Table 3.1), we also include country pair fixed effects, ω_p, in (3.2).

We also examine whether the treatment effect interacts with baseline AS characteristics by including interactive terms:

$$\mathrm{Y}_{ict} = \theta_0 + \theta_1 \mathrm{D}_c + \theta_2 \mathrm{X}_{ic} + \theta_3 \mathrm{D}_c * \mathrm{X}_{ic} + \omega_p + \varepsilon_{ict}\,, \tag{3.3}$$

where the interactive term $\mathrm{D}_c * \mathrm{X}_{ic}$ is added based on (3.2). The estimate of θ_3 captures the part of treatment effect moderated by baseline AS characteristics.

Outbound spam volumes from organizations within a country may be correlated because of common policies and regulations for a country. These country-clustered missing variables would result in highly positively correlated error terms. Failure to correct for the correlation could result in underestimated standard errors and thus overestimated treatment effects (Bertrand et al. 2004). We therefore use cluster-robust standard errors at the country level (the level of treatment assignment) in estimation of all of the above models to allow for both error heteroskedasticity and flexible within-cluster error correlation. However, the asymptotic justification based on cluster-robust standard errors assumes that the number of clusters goes to infinity. With few (five to 30) clusters, cluster-robust standard errors can still be understated (Cameron et al., 2008). Since we have only eight clusters (eight countries), this problem is likely to exist. So we further use the wild cluster bootstrap-t procedure suggested by Cameron et al. (2008) to adjust the estimated standard errors for θ_1 in our main models (3.2)).

3.6 Baseline Comparison

Since assigning firms into treatment and control groups is not random, it is necessary to test the difference in pre-treatment conditions that may be correlated with the outbound spam. If the difference is not statistically significant, then any differences in post-intervention outcomes between the two groups can be causally attributed to the intervention. Otherwise, the pre-treatment difference needs to be controlled in order to make a precise estimation on treatment effect. To check whether firm characteristics were similar or not between the two groups, we run regressions of the number of IP addresses and pre-treatment baseline spam volume (average spam volume for March and April 2012) on treatment status using (3.1).

We present the comparison of firms at baseline in Table 3.4. Column 1 contains the average characteristics for the control group. Columns 2 and 3 present the estimated differences between the treatment and control groups. The results in column 2 do not include any controls, whereas those in column 3 control for country pair fixed effects. The differences in average baseline spam and IP number are statistically significant and large in magnitude. Specifically, the organizations in the treatment group generated about

50% less spam than those in the control group before the treatment. On average, the organizations in the treatment group also have about four times more IP addresses than those in the control group. Both two variables are likely to be correlated with post-treatment outbound spam.

Table 3.4 Baseline Comparison

	Control Mean	Treatment Difference No Controls	Treatment Difference Country Pair FE
	(1)	(2)	(3)
Baseline spam	218439	-106540	-152449**
		(111622)	(57470)
IP number	140495	647773*	625859*
		(327720)	(277616)
Observations	540	1717	1717

Notes. Column 1 contains the average characteristic of the organizations in the control countries. Columns 2 and 3 contain estimates of the average difference in characteristics between the control and treatment organizations, without controls and with controls for country pair fixed effects, respectively. Standard errors are clustered by country and shown in parentheses. * indicates statistical significance at the 10% level, ** at the 5% level, and *** at the 1% level.

3.7 Information Disclosure Effect

The estimation of the effect of information disclosure is based on comparing the outgoing spam volumes of the treatment group and the control group, according to (3.1) and (3.2). The results are presented in Table 3.5. Column (1) displays the results from the basic model in (3.1), where only treatment indicator is included. It shows that although the treatment organizations sent out less spam than the control organizations, the difference is statistically insignificant. However, once we control for country pair fixed effect, the difference becomes significant and also increases in magnitude (Column (2)).

Table 3.5: Effect of Information Disclosure

	(1) Basic model	(2) Basic model +Country pair FE	(3) Basic model + Country pair FE +Controls	(4) Basic model + Country pair FE +Controls
	Spam	Spam	Spam	Ln(Spam)
Constant	121248* (52992)	163213*** (31879)	-1274 (4250)	6.4448*** (0.3441)
Treatment	-81393 (53820)	-103197** (30849)	-17757** (4076)	-2.8197*** (0.3669)
Baseline spam			0.4922*** (0.0160)	0.0000003 (0.0000002)
IP number			-0.0058*** (0.0005)	0.0000002*** (0.00000002)
Significance level using wild bootstrap-t			0.002	0.002
Observations	14255	14255	14248	14248

Notes. Column 1 displays the estimate of treatment effect on outgoing spam using the basic model without controls (3.1)). Column 2 reports the result controlling for country pair fixed effects. Column 3 controls for country pair fixed effects, baseline spam volume, and the number of IP addresses. Column 4 reports the estimate of treatment effect on log transformed outgoing spam controlling for country pair fixed effects, baseline spam volume, and the number of IP addresses. All standard errors are clustered by country and shown in parentheses. Row significance level using bootstrap-t reports the significance level for the estimate of treatment effect in Column 3. * indicates statistical significance at the 10% level, ** at the 5% level, and *** at the 1% level.

As suggested by Table 3.1, the treatment organizations significantly differ from the control organizations in terms of baseline spam and number of IP addresses. Therefore, in addition to country pair fixed effect, Column 3 also controls for baseline spam volume and number of IP addresses and contains the main results. It is not surprising that both baseline spam volume and the number of IP addresses significantly affect post-treatment outbound spam. Baseline spam volume is positively correlated with spam volume during the treatment period, indicating the persistence of certain security vulnerabilities. Unless the subject organization takes efforts to deal with these vulnerabilities, it will be continuously exploited by malicious attackers. The number of IP addresses is found to be negatively correlated with outgoing spam volume, suggesting that large systems tend to have less vulnerability. On the one hand, large systems provide attackers with more opportunities. On the other hand, large systems are likely to invest more in Internet security. Our finding suggests that the later force dominates the former.

Controlling for baseline spam and number of IP addresses drops the estimate of treatment effect by approximately 80% from 103,197 to 17,757, suggesting that the baseline spam and number of IP addresses explain a substantial part of variation in post-treatment outbound spam. Nevertheless, the treatment effect remains significant and sizable even with controls for the two characteristics. Given that the average outgoing spam volume for the treatment group is 111,899, the size of this effect is estimated at approximately 15.9%. Romanosky et al. (2011) found that the adoption of data breach disclosure laws can reduce identity theft caused by data breaches by 6.1% on average. Although the two findings are consistent, the comparison suggests that public information disclosure is more effective than notifying only those who have been affected.

Considering the small number of clusters, we use the wild cluster bootstrap-t procedure suggested by Cameron et al. (2008) to further test the treatment effect estimate. The bootstrap result shows that the estimate is robust to such asymptotic refinement. In addition, to test whether the estimate is subject to the functional specification of the statistical model, we take log transformation of spam volume, which smooths out the skewness in the distribution of spam, and run the same estimation again. According to the results presented in column (4), the treatment effect becomes even more significant statistically. Therefore, we can safely arrive at the conclusion that public disclosure of outbound spam does help reduce outbound spam.

To examine how the treatment effect dynamically changes as the treatment proceeds, we run the estimation for each month of the treatment period separately, controlling for country pair fixed effects, baseline spam, and the number of IP addresses. The results for all seven months of the treatment period are presented in Table 3.6. Throughout the entire period, the estimates of treatment effect are consistent and increase in magnitude, which provides additional support for our conclusion.

Table 3.6: Effect of Information Disclosure by Time

	(1) 1 month	(2) 2 months	(3) 3 months	(4) 4 months	(5) 5 months	(6) 6 months	(7) 7 months
Constant	7429 (6124)	-10282* (4505)	-9821 (5828)	-7308 (5405)	-2276 (8665)	-6402 (7918)	9327* (4323)
Treatment	-15476* (7296)	-10843* (5229)	-14734** (5925)	-26353*** (4257)	-36988** (9620)	-18685 (12003)	-13304** (4541)
Baseline spam	0.3049*** (0.0043)	0.5125*** (0.0140)	0.4725*** (0.0141)	0.5256*** (0.0160)	0.7260*** (0.0332)	0.7520*** (0.0342)	0.3799*** (0.0103)
IP number	0.0012 (0.0012)	-0.0036** (0.0008)	-0.0035** (0.0014)	-0.0045** (0.0019)	-0.0131** (0.0029)	-0.0149** (0.0032)	-0.0051*** (0.0006)
Observations	1717	1717	1717	1717	1717	1717	1717

Notes. Columns 1 to 7 display the estimates for the first to seventh month after the treatment. Country pair fixed effects, baseline spam, and number of IP addresses are included in all estimations. Standard errors are clustered by country and shown in parentheses. * indicates statistical significance at the 10% level, ** at the 5% level, and *** at the 1% level.

Then, we allow the treatment effect to interact with pre-treatment characteristics to see whether the effect will differ on different organizations. Table 3.7 presents the results. Column (1) simply displays the main results from column 3 of Table 3.5 without any interaction as a benchmark for comparison. Columns (2) and (3) present the results allowing the treatment effect to interact with baseline spam volume and the number of IP addresses, respectively. Column (4) presents the results allowing the treatment effect to interact with the baseline rank, which is the spam ranking among all organizations in the same country at the time one month before treatment. Column (5) presents the results allowing the treatment effect to interact with baseline top10, which is a binary indicator for whether the organization ranked top 10 among all organizations in the same country at one month before treatment. Country pair fixed effects, baseline spam, and number of IP addresses are included in all columns.

Table 3.7: Interaction Effect of Information Disclosure

	(1) No Interaction	(2) Interact with baseline spam	(3) Interact with IP number	(4) Interact with baseline rank	(5) Interact with baseline top10
Constant	-1274 (4250)	-1478 (2363)	-1746 (2945)	-49205* (23181)	-9727 (10995)
D_c	-17757** (4076)	10656 (9605)	-17091*** (2791)	-86068 (47539)	-4212 (13915)
Baseline spam	0.4922*** (0.0160)	0.5027*** (0.0006)	0.4919*** (0.0166)	0.4942*** (0.0133)	0.4944*** (0.0128)
IP number	-0.0058*** (0.0005)	-0.0002 (0.0007)	-0.0002 (0.0193)	-0.0036** (0.0014)	-0.0049*** (0.0007)
D_c* Baseline spam		-0.2524** (0.0776)			
D_c* IP number			-0.0057 (0.0191)		
Baseline rank				244.2** (89.05)	
D_c* Baseline rank				-336.6 (217.7)	
Baseline top10					7820 (90440)
D_c* Baseline top10					-294164 (294193)
Observations	14248	14248	14248	14248	14248

Notes. Column 1 displays the main results without any interaction effect from column 3 of Table 5 for comparison. Columns 2 and 3 present the results allowing the treatment effect to interact with baseline spam volume and number of IP addresses respectively. Column 4 and 5 present the results allowing the treatment effect to interact with the baseline rank and baseline top10 respectively. Country pair fixed effects, baseline spam, and number of IP addresses are included in all estimations. Standard errors are clustered by country and shown in parentheses. * indicates statistical significance at the 10% level, ** at the 5% level, and *** at the 1% level.

According to these results, the treatment effect does not interact with pre-treatment characteristics expect for baseline spam volume. The significant negative coefficient (-0.2524) for D_c* Baseline spam shows that information disclosure is more effective on organizations with more baseline spam. Although we listed only the top 20 spamming organizations, organizations currently off the list were also encouraged to take effort to remain that way. However, when public disclosure was imposed, organizations with severe outgoing spam problem had stronger incentives to deal with the problem to reduce reputation loss. This could have been partially because of the specific presentation we used in the present study, that we disclosed the worst behavior instead of advocating the best practice.

Since the data we used were collected repeatedly for the same sample for many months, the outcome may be serially correlated, and the resulting standard errors may be inconsistent. Besides bootstrap, Bertrand et al. (2004) proposed that the correction that collapses the time series information into a "pre"- and "post"-period can explicitly take into account the effective sample size. Using this method as an additional robustness check, we collapse our data into a pre-treatment and a post-treatment period by taking the average of spam volume for the months before the treatment and the months with treatment for each sample AS. Then we run the statistical model using the collapsed data. The results (in Table 3.8) are consistent with the results using original data.

Table 3.8: Information Disclosure Effect Using Average Spam Volume for Pre- and Post-Treatment Periods

	(1) Basic model	(2) Basic model +Country pair FE	(3) Basic model + Country pair FE +Controls
	Post-spam	Post-spam	Post-spam
Constant	112267* (53094)	157823*** (34448)	-1694 (4138)
Treatment	-72612 (53993)	-96041** (32421)	-17333*** (4261)
Pre-spam			0.4928*** (0.0158)
IP number			-0.0058*** (0.0006)
Observations	1718	1718	1718

Notes. Standard errors are clustered by country and shown in parentheses. * indicates statistical significance at the 10% level, ** at the 5% level, and *** at the 1% level.

3.8 Social Comparison Effect

As we showed in the previous section, information disclosure does positively encourage organizations to reduce outgoing spam. Although we release information only on the top-most spamming organizations, the disclosure has a positive spillover effect on organizations that are currently off the list. The disclosure effect mainly comes from the concern for reputation loss; that is, the more spam an organization sends out, the more likely it will be listed on SpamRankings.net. Another force that would potentially affect organizations' behavior is the social comparison effect caused by the information of other organizations' behavior. The important information on others' behavior released on our website is the maximum spam volume observed from the most spamming AS (referred to as Max spam) and the minimum listed spam volume observed from the AS that ranked 10th in the treatment country (referred to as Min spam). According to social comparison theory (Festinger 1954), organizations will react to the specific information disclosed as well as to the disclosure mechanism in that organizations have the tendency to behave in consistency with others. Therefore, in addition to the treatment variable, we add two variables, Treatment*Max spam and Treatment*Min spam, to test whether treatment organizations will react to the specific information disclosed on other organizations.

Table 3.9 Impact of Specific Information Disclosed

	(1) Basic model + Country pair FE +Controls	(2) Basic model + Country pair FE +Controls	(3) Basic model + Country pair FE +Controls
	Spam	Spam	Spam
Constant	353.3 (5172)	-44.00 (5041)	375.0 (5271)
Treatment	-24601*** (6320)	-17514** (7266)	-24922** (7375)
Baseline spam	0.4899*** (0.0177)	0.4899*** (0.0177)	0.4899*** (0.0177)
IP number	-0.0066*** (0.0014)	-0.0066*** (0.0014)	-0.0066*** (0.0014)
Treatment*Max spam	0.0002* (0.0001)		0.0002** (0.00008)
Treatment*Min spam		-0.0051 (0.0114)	0.0007 (0.0096)
Observations	14248	14248	14248

Notes. Standard errors are clustered by country and shown in parentheses. * indicates statistical significance at the 10% level, ** at the 5% level, and *** at the 1% level.

Table 3.9 shows the results after including these two additional variables. The coefficients indicate how the specific information would modify the treatment effect. We find that organizations react only to the specific Max spam, not the Min spam, meaning that they pay attention only to the worst spam sending behavior but not to the mediocre behavior. The coefficient estimate of Treatment*Max spam is positive, suggesting that the more outgoing spam observed from the worst behavior, the less likely organizations will be to take effort to improve their own behavior. In other words, it shows that if even the worst behavior is not so bad, organizations will have more pressure or desire to improve themselves. This finding is consistent with the prediction by social comparison process, during which individuals evaluate themselves against others and the existence of a discrepancy leads to actions toward reducing it.

3.9 Concluding Remarks

Governments, businesses, and consumers are constantly exposed to the risk of cybercrime. Our society has recognized the need for additional laws and co-operation to protect consumer privacy, enterprise assets, intellectual property, and critical national infrastructure. In the thriving and fast-moving discipline of Internet security, many are searching for technical solutions such as firewall and antivirus software. We propose that Internet security needs to be improved from the perspective of fundamental motivations. Systems are prone to failure when the person guarding them is not the person who suffers when they fail (Anderson and Moore 2006). An organization's security vulnerabilities are also shared by other organizations but are often kept private. The negative externality gives Internet security the feature of partial public good. The private provision of public goods often results in underinvestment because of the lack of incentives. In social psychology, the underinvestment problem is often addressed through making relevant social information available and soliciting social comparison process.

Drawing upon social comparison theory, we propose a social information provision to encourage organizations to improve their Internet security. The information disclosure mechanism can incur concern for reputation loss, whereas the specific information disclosed can incur a social comparison. Through making information on other organizations' behavior publicly available, we aimed to solicit the comparison among them and impose reputation loss on those who do not behave pro-socially. To dose up the comparison, we disclosed the relative rankings for all of the organizations in a country in addition to absolute performance and listed the top 10 worst organizations by their standings. Such "shame" lists make bad behaviors notorious. Using a field quasi-experiment on outgoing spam for 1,718 ASes in eight countries, we show that providing social information on outgoing spam encouraged the treatment organizations to reduce it by approximately 15.9%. As compared with an existing study (Romanosky et al. 2011), which documented a 6.1% effect of adopting data breach disclosure laws on identity theft, this result shows that making social information publicly available is more effective than notifying only affected consumers in motivating desirable pro-social behavior.

We find a positive spillover effect in that even though only the top 10 worst organizations are listed, such listing incentivizes both listed and unlisted organizations. However, the impact is stronger on organizations with more spam. The number of IP addresses is found to be negatively correlated with outgoing spam volume. The number of IP addresses measures the size of the AS. On the one hand, large ASes have more incentives to invest in Internet security because of economies of scale. On the other hand, large ASes have more exploitable opportunities for hackers. The result indicates the former dominates the latter. A closer look at the social comparison process reveals that the more outgoing spam observed from the worst spammer, the less likely an organization will be to improve its own behavior. This finding again reflects that improving Internet security requires collective work of all organizations and that individual behavior can generate strong positive externalities on others. We can utilize the desire for information on others' behavior and use of the information for self-evaluation to intervene in individual organizations' security decisions and solicit the desirable behavior.

Our present study has implications for information architecture design and public policy making. With the ubiquity of Internet, the things that people do online can be tracked, providing an abundance of data. The question we currently face is not the lack of data but how to make use of the data available. Online users care about their popularity, reputation, and social status within the community. If we can capture users' actions, aggregate and display the relevant information, and provide the right feedback as the right intervention at the right time, we can lead their behaviors in our intended direction. With respect to public policy, our present work is among the few empirical studies on Internet security using security vulnerabilities data. Policy makers have hesitated to use security information disclosure for a long time. Although a fierce argument has been observed surrounding disclosure, little attention has been paid to information display or presentation. We believe what is more important than disclosure is whether the information is easy for users to interpret and compare. In the present study, we used relative rankings to enhance the disclosure effect. For policy evaluation, more information presentation methods can be considered and compared before carrying out the policy extensively. Field experimentation provides an efficient and effective method to evaluate potential policies beforehand. The same approach applies to other security, social, or environmental problems such as energy conservation and pollution. In the case where data is not available, the legislation that requires mandatory reporting can be employed to collect data.

Our present article is only our first step in studying Internet security and relevant public policy issues from social psychology and economics perspectives. We are planning to further extend the present study in several dimensions. First, we experimented only with ranking information in our study to focus on relative standing. To identify the exact effect of using ranking information versus absolute volume information, a new treatment group can be added by which organizations receive information only on absolute outbound spam volume with organizations listed alphabetically. With the established visibility of SpamRankings.net, we can experiment with more countries, more industries, and more treatment conditions. Second, the observation of reduction in outgoing spam may or may not reflect the improvement in overall Internet security. If overall Internet security improves while spam decreases, it indicates that companies take the initiative to improve their overall information security, affecting both vulnerability to spam and other threats such as phishing. This would mean that broad improvements in information security can be achieved by presenting public information on certain security issues. It is also possible that in response to public information disclosure of outbound spam, organizations may take effort to address only outbound spam issue but will still ignore other security problems. If this happens, it means that companies instead need to be individually incentivized to make improvements on individual dimensions of security. As a result, we can encourage companies to make anti-spam improvements by releasing social information on spam. However, to encourage companies to prevent phishing, we need to also release phishing information. With phishing data in addition to spam data, we can distinguish these two possibilities by exploring their correlations. In addition, we can drill down outgoing spam to botnets or snowshoe spammers to consider attackers' reactions to information disclosure.

Appendix: Proof of Results in Chapter 2

Proof of Proposition 2.1

Proof. From (2.2), we can obtain:

$$\Pi_{T+1} - \Pi_T$$
$$= \Phi\left(\frac{\alpha_{T+1}X_j'\beta + (1-\alpha_{T+1})V_j - c_j}{g_{T+1}}\right) - \Phi\left(\frac{\alpha_T X_j'\beta + (1-\alpha_T)V_j - c_j}{g_T}\right).$$

If $V_j - X_j'\beta > \frac{1}{\xi_T}\left(\frac{1}{g_T} - \frac{1}{g_{T+1}}\right)\left(X_j'\beta - c_j\right)$, where $\xi_T = \frac{1-\alpha_{T+1}}{g_{T+1}} - \frac{1-\alpha_T}{g_T} > 0$, we have $\Pi_{T+1} - \Pi_T > 0$. Thus, if a positive surprise is sufficiently large, then Π_T is increasing in T. Similarly, we can show that if a negative surprise is sufficiently large, then Π_T is decreasing in T.

Proof of Proposition 2.2

Proof. From (2.2), we can obtain:

$$\frac{\partial}{\partial \rho_{V_j}}(\Pi_{T+1} - \Pi_T)$$

$$= \phi\left[\frac{\alpha_{T+1}X_j'\beta + (1-\alpha_{T+1})V_j - c_j}{\sqrt{\frac{\left(\rho_{\varepsilon_j} + \varsigma_{T+1}\right)}{\left(\rho_{V_j} + \rho_{\varepsilon_j} + \varsigma_{T+1}\right)^2}}}\right]\frac{X_j'\beta}{\sqrt{\rho_{\varepsilon_j} + \varsigma_{T+1}}}$$

$$- \phi\left[\frac{\alpha_T X_j'\beta + (1-\alpha_T)V_j - c_j}{\sqrt{\frac{\left(\rho_{\varepsilon_j} + \varsigma_T\right)}{\left(\rho_{V_j} + \rho_{\varepsilon_j} + \varsigma_T\right)^2}}}\right]\frac{X_j'\beta}{\sqrt{\rho_{\varepsilon_j} + \varsigma_T}}$$

where $\varsigma_T = \sum_{t=2}^{T} d_{it}$, and $\phi(\cdot)$ is the density of a standard normal distribution. Because the positive surprise $(V_j - X_j'\beta)$ is sufficiently large, from Proposition 2.1, we have $\Pi_{T+1} - \Pi_T > 0$, and

$$\frac{\alpha_{T+1}X_j'\beta + (1-\alpha_{T+1})V_j - c_j}{\sqrt{\left(\rho_{\varepsilon_j} + \varsigma_{T+1}\right)/\left(\rho_{V_j} + \rho_{\varepsilon_j} + \varsigma_{T+1}\right)^2}} > \frac{\alpha_T X_j'\beta + (1-\alpha_T)V_j - c_j}{\sqrt{\left(\rho_{\varepsilon_j} + \varsigma_T\right)/\left(\rho_{V_j} + \rho_{\varepsilon_j} + \varsigma_T\right)^2}}$$

$$= \frac{X_j'\beta + (1-\alpha_T)\left(V_j - X_j'\beta\right) - c_j}{\sqrt{\left(\rho_{\varepsilon_j} + \varsigma_T\right)/\left(\rho_{V_j} + \rho_{\varepsilon_j} + \varsigma_T\right)^2}} > 0,$$

where the inequality follows from the sufficiently large positive surprise, $(V_j - X_j'\beta)$.

Thus, we have:

$$\phi\left[\frac{\alpha_{T+1}X_j'\beta + (1-\alpha_{T+1})V_j - c_j}{\sqrt{\left(\rho_{\varepsilon_j} + \varsigma_{T+1}\right) / \left(\rho_{V_j} + \rho_{\varepsilon_j} + \varsigma_{T+1}\right)^2}}\right] < \phi\left[\frac{\alpha_T X_j'\beta + (1-\alpha_T)V_j - c_j}{\sqrt{\left(\rho_{\varepsilon_j} + \varsigma_T\right) / \left(\rho_{V_j} + \rho_{\varepsilon_j} + \varsigma_T\right)^2}}\right].$$

As a result, when the positive surprise is sufficiently large, we can show that $\frac{\partial}{\partial \rho_{V_j}}(\Pi_{T+1} - \Pi_T) < 0$. Part (b) can be proved similarly.

Bibliography

[1] Aguirreagabiria, V., and Mira, P. Sequential estimation of dynamic discrete games. *Econometrica*, 75, (2007), 1–53.

[2] Aguirreagabiria, V., and Mira, P. Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica*, 70, (2002), 1519–1543.

[3] Akerlof, G. A. A theory of social custom, of which unemployment may be one consequence. *The Quarterly Journal of Economics,* 94, 4 (1980), 749-775.

[4] Amemiya, T. Regression analysis when the dependent variable is truncated normal.*Econometrica*, 41, 6 (1973), 997–1016.

[5] Anderson, R. Why information security is hard: An economic perspective. Proceedings of the 17th Annual Computer Security Applications Conference (ACSAC), 2002, New Orleans, LA.

[6] Anderson, R., and Moore, T. The economics of information security. *Science*, 27, 314 (2006), 610-613.

[7] Anderson, S.P., and Gans, J.S. Platform siphoning: Ad-avoidance and media content. *American Economic Journal*: *Microeconomics*, 3, 4 (2011), 1–34.

[8] Aral, S., Muchnik, L., and Sundararajan, A. Distinguish influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106, 51 (2009), 21544-21549.

[9] Aral, S., and Walker, D. Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science*, 57, 9 (2011), 1623-1639.

[10] Arcidiacono, P., and Miller, R.A. Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79, 6 (2011), 1823–1867.

[11] Arellano, M., Discrete choices with panel data. *Investigaciones Economica*, 25, (2000), 423–458.

[12] Arora, A., Krishnan, R., Nandkumar, A., Telang, R., and Yang, Y. Impact of vulnerability disclosure and patch availability: An empirical analysis. *Third Workshop on the Economics of Information Security*, 2004a, 1268-1287.

[13] Arora, A., Telang, R., and Xu, H.. Timing disclosure of software vulnerability for optimal social welfare. *Proceedings of the 3rd Workshop of Economic Information Systems*, 2004b, 1–47.

[14] Arthur, J.B., and Jelf, G.S. The effects of gainsharing on grievance rates and absenteeism over time. *Journal of Labor Research*, 20, 1 (1999), 133–145.

[15] Atkinson, S.; Stanley, L.; and Tschirhart, J. Revenue sharing as an incentive in an agency problem: An example from the National Football League. *RAND Journal of Economics*, 19, 1 (1988), 27–43.

[16] Bajari, P.B.; Benkard, C.L.; and Levin, J. Estimating dynamic models of imperfect competition. *Econometrica*, 75, 5 (2007), 1331–1370.

[17] Banerjee, A. V. A simple model of herd behavior. *Quarterly Journal of Economics*, 107, 3 (1992), 797-817.

[18] Bass, E.M. A new product growth for model consumer durables. *Management Science*, 15, 5 (1969), 215–227.

[19] Bateman, P.J.; Gray, P.H.; and Butler, B.S. The impact of community commitment on participation in online communities. *Informational Systems Research*, 22(2011), 841-854.

[20] Bernheim, B. D. A theory of conformity. Journal of Political Economy, 102, 5 (1994), 841-877.

[21] Berry, S.; Levinsohn, J.; Pakes, A. Automobile prices in market equilibrium. *Econometrica*, 63, 4(1995), 841–890.

[22] Bertrand, M., Duflo, E., and Mullainathan, S. How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119, 1 (2004), 249-275.

[23] Black, S.E., and Lynch, L.M. What's driving the new economy? The benefits of workplace innovation. *Economic Journal*, 114, 493 (2004), 97–116.

[24] Blundell, R.; Browning, M.; and Crawford, I. Nonparametric engel curves and revealed preference. *Econometrica*, 71, 1(2003), 205–240.

[25] Butler, B. S. Membership size, communication activity, and sustainability: A resource-based model of online social structures. *Information Systems Research*, 12, 4 (2001), 346-362.

[26] Cai, H., Chen, P., and Fang, H. Observational learning: evidence from a randomized natural field experiment. *American Economic Review*, 99, 3 (2009), 864-882.

[27] Cain, D.M., Loewenstein, G., and Moore, D. A. The dirt of coming clean: Perverse effects of disclosing conflicts of interest. *Journal of Legal Studies*, 34, (2005), 1-25.

[28] Cameron, A. C., Gelbach, J. B., and Miller, D. L. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90, 3 (2008), 414–427.

[29] Campbell, K., Gordon, L. A., Loeb, M. P., and Zhou, L. The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *Journal of Computer Security*, 11, (2003), 431–448.

[30] Cate, F. Comparative approaches to security breaches. Symposium on Security Breach Notification Six Years Later: Lessons Learned about Identity Theft and Directions for the Future, Berkeley Center for Law and Technology, Berkeley, CA, 2009.

[31] Chen, Y., Harper, F. M., Konstan, J., and Li, S. X. Social comparisons and contributions to online communities: A field experiment on MovieLens. *The American Economic Review*, 100, 4 (2010), 1358-1398.

[32] Chevalier, J., and Mayzlin, D. The effect of word of mouth on sales: Online book review. *Journal of Marketing Research*, 43, 3 (2006), 345-354.

[33] D'Arcy, J., Hovav, A., and Galletta, D. User awareness of security countermeasures and its impact on information systems misuse: A deterrence approach. *Information Systems Research*, 20, 1 (2009), 79-98.

[34] Dahlman, C. J. The problem of externality. *Journal of law and economics*, 22, 1 (1979), 141-162.

[35] Dellarocas, C. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49, 10 (2003), 1407–1424.

[36] Dellarocas C.; Gao G.; and Narayan, R. Are consumers more likely to contribute online reviews for hit products or niche products? *Journal of Management Information Systems*, 27, 2 (Fall 2010), 127–158.

[37] Dillenberger D., and Sadowski, P. Ashamed to be selfish. *Theoretical Economics*, 7, 1 (2012), 99-124.

[38] Doraszelski, U., and Jaumandreu, J. R&D and productivity: Estimating production functions when productivity is endogenous. Manuscript, Department of Economics, Harvard University, 2008.

[39] Duan, W., Gu, B., and Whinston, A.B. The dynamics of online word-of-mouth and product sales – An empirical investigation of the movie industry. *Journal of Retailing*, 84, 2 (2009), 233-242.

[40] Duflo, E., Dupas, P., and Kremer, M. Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in Kenya. *American Economic Review*, 101, 5 (2011), 1739-1774.

[41] Eliaz, K., and Spiegler, R. On the strategic use of attention grabbers. *Theoretical Economics*, 6, 1 (2001), 127-155.

[42] Fang, Y., and Neufeld, D. Understanding sustained participation in open source software projects. *Journal of Management Information Systems*, 25, 4 (Spring 2009), 9–50.

[43] Festinger, L. A theory of social comparison processes. *Human Relations*, 7, (1954), 117 -140.

[44] Frei, S., The security of end-user PCs: An empirical analysis. DDCSW: Collaborative Data-Driven Security for High Performance Networks, Internet2 and WUSTL, August, 2010.

[45] Frey, B. S., and Meier, S. Social comparisons and pro-social behavior – testing "conditional cooperation" in a field experiment. *American Economic Review*, 94, (2004), 1717–1722.

[46] Fung, A., Graham, M., and Weil, D. Full Disclosure: The Perils of and Promise of Transparency. Cambridge, MA: Cambridge University Press, 2007.

[47] Gal-Or, E., and Ghose, A. The economic incentives for sharing security information. *Information Systems Research*, 16, 2 (2005), 186-208.

[48] Ghose, A., Han, S., and Iyengar, R. An empirical analysis of network characteristics and social contagion on the mobile internet. *Working Paper*, (2012).

[49] Godes, D., and Mayzlin, D. Firm-created word-of-mouth communication: Evidence from a field test. *Marketing Science*, 28, 4 (2009), 721-739.

[50] Goldenberg, J., Han, S., Lehmann, D.R., and Hong, J.W. The role of hubs in the adoption process. *Journal of Marketing*, 73, 2(2009), 1-13.

[51] Goldenberg, J., Libai, B., and Muller, E. The chilling effects of network externalities. *International Journal of Research in Marketing*, 27, 1(2010), 4-15.

[52] Goldenberg, J., Oestreicher-Singer, G., and Reichman, S. The quest for content: The integration of product networks and social networks in online content exploration. *Journal of Marketing Research*, 49, 4(2012), 452-468.

[53] Goodman, J., Cormack, G. V., and Heckerman, D. Spam and the ongoing battle for the inbox. *Communications of the ACM*, 50, 2 (2007), 24–33.

[54] Gordon, L. A., Loeb, M., and Lucyshyn, W. Sharing information on computer systems security: An economic analysis. *Journal of Accounting Public Policy*, 22, 6 (2003), 461–485.

[55] Griliches, Z. R&D, education, and productivity: A retrospective. Harvard University Press, Cambridge, 2000.

[56] Griskevicius, V., Tybur, J. M., and Van den Bergh, B. Going green to be seen: Status, reputation, and conspicuous conservation. *Journal of Personality and Social Psychology*, 98, 3 (2010), 392-404.

[57] Hendricks, K., and Sorensen, A. Information and the skewness of music sales. *Journal of Political Economy*, 117, 2 (2009), 324–369.

[58] Ho, T., Li, S., Park, S.E., and Shen, Z.M. Customer influence value and purchase acceleration in new product diffusion. *Marketing Science*, 31, 2(2012), 236-256.

[59] Holmstrom, B. Managerial incentive problems: A dynamic perspective. *Review of Economic Studies*, 66, 1 (1999), 169–182.

[60] Hopkins, E., and Kornienko, T. Running to keep in the same place: Consumer choice as a game of status. *American Economic Review*, 94, 4 (2004), 1085-1107.

[61] Horner, J. Reputation and competition. *The American Economic Review*, 92, 3 (2002), 644–663.

[62] Hotz, V. J., and Miller, R. A. Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, 60, 3 (1993), 497–529.

[63] Hui, K. L., Teo, H. H., and Lee, S. Y. T. The value of privacy assurance: An exploratory field experiment. *MIS Quarterly*, 31, 1 (2007), 19-33.

[64] Isacenkova, J., and Balzarotti, D. Measurement and evaluation of a real world deployment of a challenge-response spam filter. *Proceedings of ACM ICM*, 2011.

[65] Iyengar, R., Van den Bulte, C., and Choi, J. Distinguishing among mechanisms of social contagion in new product adoption: framework and illustration. *Working Paper*, (2012).

[66] Iyengar, R., Van den Bulte, C., and Valente, T.W. Opinion leadership and social contagion in new product diffusion. *Marketing Science*, 30, 2(2011), 195-212.

[67] Jin, G. Z., and Leslie, P. The effect of information on product quality: Evidence from restaurant hygiene grade cards. *Quarterly Journal of Economics*, 118, (2003) 409–451.

[68] Jones, C.; Hesterly, W.S.; and Borgatti, S.P. A general theory of network governance: Exchange conditions and social mechanisms. *Academy of Management Review*, 22, 4 (1997), 911–945.

[69] Kankanhalli, A., Teo, H. H., Tan, B. C. Y., and Wei, K.-K. An integrative study of information systems security effectiveness. *International Journal of Information Management*, 23, 2 (2003) 139–154.

[70] Kannan, K., Rees, J., and Sridha, S. Market reactions to information security breach announcements: An empirical analysis. *International Journal of Electronic Commerce*, 12, 1 (2007), 69-91.

[71] Karau, S. J., and Williams, K. D. Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality and Social Psychology*, 65, 4 (1993), 681-706.

[72] Kraut, R. E., Morris, J., Telang, R., Filer, D., Cronin, M., and Sunder, S. Markets for attention: Will postage for email help? *CSCW '02 Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 2002, 206–215, New York, NY.

[73] Krupka, E., and Weber, R. A. The focusing and informational effects of norms on pro-social behavior. *Journal of Economic Psychology*, 30, 3 (2009), 307-320.

[74] Lakhani, K., and Hippel E. How open source software works: "Free" user-to-user assistance. *Research Policy*, 32, (2003), 923–943.

[75] Lerner, J., and Tirole, J. Some simple economics of open source. *Journal of Industrial Economics*, 50, 2 (2002), 197–234.

[76] Liu-Thompkins, Y., and Rogerson, M. Rising to stardom: An empirical investigation of the diffusion of user-generated content. *Journal of Interactive Marketing*, 26, 2(2012), 71-82.

[77] Loder, T., Van Alstyne M., and Wash, R. An economic answer to unsolicited communication. *Proceedings of the 5th ACM conference on Electronic commerce*, 2004, 40-50.

[78] Luttmer, E. F. P. Neighbors as negatives: Relative earnings and wellbeing. *The Quarterly Journal of Economics*, 120, 3 (2005), 963–1002.

[79] Ma, M., and Agarwal, R. Through a glass darkly: Information technology design, identity verification, and knowledge contribution in online communities. *Informational Systems Research*, 18, 1 (2007), 42–67.

[80] Manski, C.F. Identification of endogenous social effects: The reflection problem. *Review of Economic Studies*, 60, 3(1993), 531-542.

[81] Manski, C.F. Economic analysis of social interactions. *The Journal of Economic Perspectives*, 14, 3 (2000), 115–136.

[82] Miller, C.C. Actors in smaller studios, making pictures for the smaller screen. New York Times, 2011, (accessed April 10, 2011), [available at http://www.nytimes.com/2011/04/11/business/media/11youtube.html?_r=1&pagewanted=all].

[83] Mookerjee, V., Mookerjee, R., Bensoussan, A., and Yu, W. T. When hackers talk: Managing information security under variable attack rates and knowledge dissemination. *Information Systems Research*, 22, 3 (2011), 606-623.

[84] Moretti, E. Social learning and peer effects in consumption: Evidence from movie sales. *Review of Economic Studies*, 78, 1(2011), 356-393.

[85] Mortimer, J. Vertical contracts in the video rental industry. *The Review of Economic Studies*, 75, 1 (2008), 165–199.

[86] Moyer, E. Breach exposes chase, capital one, TiVo customers. CNET News, 2011, April 2. http://news.cnet.com/8301-1009_3-20050068-83/breach-exposes-chase-capital-one-tivo-customers/.

[87] Pakes, A.; Ostrovsky, M.; and Berry, S. Simple estimators for the parameters of discrete dynamic games. *The Rand Journal of Economics*, 38, 2 (2007), 373–399.

[88] Pigou, A. The Economics of Welfare. McMillan&Co., London, 1920.

[89] Qiu, L., and Whinston, A.B. Naive herding in location-based networks: Induced social learning and optimal dynamic pricing. *INFORMS Conference on Information Systems and Technology (CIST)*, (2012) Phoenix, AZ.

[90] Ransbotham, S., and Mitra, S. Choice and chance: A conceptual model of paths to information security compromise. *Information Systems Research*, 20, 1 (2009), 121-139.

[91] Rao, J. M., and Reiley, D. H. The economics of spam. *The Journal of Economic Perspectives*, 26, 3 (2012), 87-110.

[92] Reiss, P.C., and Wolak, F.A. Structural econometric modeling: Rationales and examples from Industrial Organization. In J.J. Heckman and E.E. Leamer (eds.), Handbook of Econometrics 6, Amsterdam: Elsevier Science, North-Holland, 2005, pp. 4277–4415.

[93] Rice, S.C. Reputation and uncertainty in online markets: An experimental study. *Informational Systems Research*, 23, (2012), 436-452.

[94] Romanosky, S., Telang, R., and Acquisti, A. Do data breach disclosure laws reduce identify theft? *Journal of Policy Analysis and Management*, 30, 2(2011), 256-286.

[95] Rust, J. Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, 55, 5 (1987), 999–1033.

[96] Rust, J. Structural estimation of Markov decision processes. Handbook of Econometrics, 4, (1994), 3081–3143.

[97] Saito, K. Role conflict and choice: Shame, temptation, and justifications. *Working paper*, (2011).

[98] Schneier, B. Computer security: It's the economics, stupid. Workshop on Economics and Information Security, 2002, University of California, Berkeley, CA.

[99] Schwartz, P., and Janger, E. Notification of data security breaches. *Michigan Law Review*, 105(2007), 913-984.

[100] Shang, J., and Croson, R. A field experiment in charitable contribution: The impact of social information on the voluntary provision of public goods. *The Economic Journal*, 119, 540(2009), 1422-1439.

[101] Short, J., Williams, E., and Christie, B. The Social Psychology of Telecommunications. Hoboken, NJ: Wiley, 1976.

[102] Sipior, J. C., Ward, B. T., and Bonner P. G. Should spam be on the menu? *Communications of the ACM*, 47, 6(2004), 59-63.

[103] Stewart, D. Social status in an open source software community. *American Sociological Review*, 70, 5 (2005), 823–842.

[104] Straub, D. W., and Welke, R. J. Coping with systems risk: Security planning models for management decision making. *MIS Quarterly*, 22, 4(1998), 441–469.

[105] Susarla, A.; Oh, J.; and Tan, Y. Social networks and the diffusion of user-generated content: Evidence from YouTube. *Informational Systems Research*, 23, 1(2012), 23-41.

[106] Van Den Bulte, C., and Stremersch, S. Social contagion and income heterogeneity in new product diffusion: A meta-analytic test. *Marketing Science*, 23, 4(2004) 530-544.

[107] Wasko, M.M., and Faraj, S. Why should I share? Examining knowledge contribution in networks of practice. *MIS Quarterly*, 29, 1 (2005), 35–57.

[108] Wilbur, K. C. A two-sided empirical model of television advertising and viewing markets. *Marketing Science*, 27, 3(2008), 356-378.

[109] Wooldridge, J. M. Econometric Analysis of Cross Section and Panel Data. Cambridge, MA: MIT Press, 2002.

[110] Wooldridge, J. M. What's new in Econometrics? Lecture 10 difference-in-differences estimation," NBER Summer Institute, mimeo, (2007).

[111] Yue, W., and Cakanyildirim, M. Intrusion prevention in information systems: Reactive and proactive response. *Journal of Management Information Systems*, 24, 1(2007), 329–353.

[112] Zhang, X., and Zhu, F. Group size and incentives to contribute: A natural experiment at Chinese Wikipedia. *American Economic Review*, 101, 4 (2011), 1601–1615.

Vita

Qian Tang received a bachelor's degree in Information Management and Information Systems from Tsinghua University in 2006 and a master's degree in Management Science and Engineering from Tsinghua University in 2008. In August 2008, she was admitted to the graduate program at the University of Texas at Austin. In 2012, she received the master of science degree in Economics at the University of Texas at Austin.

Permanent email: qiantang@utexas.edu

This dissertation was typed by the author.