Understanding Online Two-sided Market: An Empirical Perspective

Jinyang Zheng

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2017

Reading Committee:

Yong Tan, Chair

Ming Fan

Hyeunjung (Elina) Hwang

Program Authorized to Offer Degree:

Business Administration

University of Washington

Abstract

Understanding Online Two-sided Market: An Empirical Perspective

Jinyang Zheng

Chair of the Supervisory Committee:
Neal and Jan Dempsey Professor Yong Tan
Information Systems and Operations Management at Foster School of Business

This dissertation provides a comprehensive investigation of online two-sided markets from empirical perspective. The four essays included propose research questions from four different angles and answer those questions by using four correspondingly matched empirical setups. In essay 1, we study the adoption of online two-sided market by investigating how two-sided sales promotion affects drivers' willingness to use the TNC app and how the TNC develops its optimal promotion strategies accordingly. We find that the substantial value of early promotion not only encourages current usage but also fosters learning that sustains drivers' use of the app afterward. The results also show that revealed tips from passengers signal low quality of service and that platform cashback to passengers has a positive effect on drivers by increasing drivers' chances of

being rewarded. In essay 2, we investigate the information asymmetry problem of online two-sided market by measuring the effect of platform endorsements and consumer-generated reputation on demand in the online service marketplace and the effect of a "conform or be cast out" policy that is applied by the platform to ensure that sellers have platform refund insurance on the platform-wide quality and competency. We find that individuals have relatively consistent sensitivity to consumer-generated reputation, but different perceptions of platform endorsements, some of which could have a negative effect on demand. Regarding the "conform or be cast out" policy, we find that, even though casting-out reduces the variety of sellers and, thus, decreases platform-wide demand and consumer welfare, the negative effect is offset by sellers' having platform refund insurance. In essay 3, we study the social influence of online two-sided market by investigating how a review-in-review (RIR) affects rating behavior. We find measurable evidence that individuals rate, in part, to satisfy their expectation of gaining social capital. The average rater weights social capital gain at approximately 20%. In essay 4, we examined the role of the mega player in the online two-sided market by estimating the spillover effects of WeChat on the other top-50 most frequently used apps in China. We find that the spillover effects of WeChat are limited; in fact, only two other only two other apps, Tencent News and Taobao, receive positive spillover effects from WeChat.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to express my sincere appreciations to my advisor, Prof. Yong Tan, for the continuous support of my Ph.D. study and related research, for his understanding, patience, motivation, and immense knowledge. I could not have imagined having a better advisor and mentor for my Ph.D. study. His unconditional and endless caring makes my Ph.D. study one of the most enjoyable life experience. His enthusiasm and vision for research has greatly influenced me and will keep on motivating me to conduct high-quality research.

My thanks also go to the members of my dissertation committee: Prof. Ming Fan, Prof. Hyeunjung (Elina) Hwang, Prof. Abhishek Borah and Prof. Zaozao Zhang. Their insightful comments and encouragement help me to improve this dissertation.

I would like to extend my gratitude to all fellow doctoral students and staffs at the Foster School of Business for their long-standing help and support for the four-years study.

I am also grateful to Ms. Dorrit Bern and Mr. John R. McCabe for your generous financial support. Your support and help increase my chances of success in my research career.

DEDICATION

I dedicate my dissertation work to my family and many friends. A special feeling of gratitude to my loving parents, Shunqing Zheng and Cailing Chen, whose protection, caring and guidance lead me to what I accomplished today.

I also dedicate this dissertation to my many friends and teachers in my high school and college. Thanks for showing me the example of being a better person and sharing me with the happiness of staying with the smart people, which motives me to keep on surpass whom I might have been.

A special dedication to my alma mater, Fudan University, the place who endows me with the philosophy as a life-being and now is the harbor that refuel smyself for the next journey.

I dedicate this work and give special thanks to my best friend, Chen Qian, who has never left my side during the toughest three years in my Ph.D. study.

Finally, for my grandma, Lianxiang He, in memoriam. You are with me and will watch over me wherever I go.

Chapter 1. INTRODUCTION

The mobile app of a transportation network company (TNC) allows the TNC platform to run aggressive and diverse two-sided sales promotions to help to introduce new products. In the first essay, we examine how two-sided sales promotion affects drivers' willingness to use the TNC app and how the TNC develops its optimal promotion strategies accordingly. To investigate the effects of sales promotion, we estimate a structural model of drivers' decisions to accept orders and to cancel generated orders and their perception of passengers' willingness to utilize a sales promotion. Bayesian learning processes are introduced to account for decisions under uncertainty as the app is introduced. We find measurable evidence of drivers' learning about the value of the attributes of the transportation network app. The results indicate that the substantial value of early promotion not only encourages current usage but also fosters learning that sustains drivers' use of the app afterward. Our results also show that revealed tips from passengers signal low quality of service and that platform cashback to passengers has a positive effect on drivers by increasing drivers' chances of being rewarded. Given the estimated parameters, we run simulations to explicitly measure the indirect effects of the sales promotion, as supported by learning, and show how cashback for passengers affects the decisions of drivers. Finally, our experimental promotion policies show improved performance with regard to drivers' willingness to use the app as well as its cost effectiveness.

The online service marketplace faces information asymmetry due to the intangibility and variability of services. To alleviate this problem, the platform utilizes a reputation system that provides standardized platform endorsements and information on the purchase history of consumers. The second essay measures the effect of platform endorsements and consumer-generated reputation on demand in the online service marketplace. In particular, we apply a demand estimation method that uses aggregate data to understand heterogeneous sensitivity to the characteristics of sellers. We further conduct a simulation study to investigate the effect of a "conform or be cast out" policy that is applied by the platform to ensure that sellers have platform refund insurance, which improves platform-wide quality and competency. Our findings show that individuals have relatively consistent sensitivity to consumer-generated reputation, but different perceptions of platform endorsements, some of which could have a negative effect on demand. With regard to the "conform or be cast out" policy, we find that, even though casting-out reduces the variety of sellers and, thus, decreases platform-wide demand and consumer welfare, the negative effect is offset by sellers' having platform refund insurance. Further, this policy is found to lead sellers into further quality escalation until a new equilibrium, which indirectly benefits platform demand and consumer welfare, is established. The study provides comprehensive analyses and insights for stakeholders in the service online marketplace, including sellers, consumers, and the platform operators.

A review-in-review (RIR) is a function that allows content consumers to generate positive or negative comments and evaluations for a primary review of a product. The introduction of an RIR has the potential to reshape the primary content (rating)-generation process because raters who seek increased social capital might be driven to rate content at a level that is expected to generate more "helpful" and fewer "unhelpful" responses. In the third essay, to investigate how an RIR affects rating behavior, we quantify the expectation-formation of the RIR by mimicking an individual's learning from observing past content and develop a partially ordinal discrete choice model that allows rating responses to reflect a mixture of ordinal latent utility, such as perceived quality, and a conditional multinomial utility, such as expected social capital gain. We find measurable evidence that individuals rate, in part, to satisfy their expectation of gaining social capital. Our results also show that there is heterogeneity in individuals' weighting of expected social capital gain and perceived quality. The average rater weights social capital gain at approximately 20%, suggesting a role of an expected RIR in rating behavior. We further experiment with the settings in which RIR features are switched off to quantitatively measure the impact of the RIR on individuals' rating decisions and their consequences for ratings distribution.

WeChat, an instant messaging app, is considered a mega app due to its dominance in terms of usage among Chinese smartphone users. Nevertheless, little is known about it externality in regard to the broader app market. The fourth essay estimates the spillover effects of WeChat on the other top-50 most frequently used apps in China through data on users' weekly app usage. Given the challenge of determining causal inference from observational data, we apply a graphical model and econometrics to estimates the spillover effects through two steps: (1) we determine the causal structure by estimating a partially ancestral diagram, using a Fast Causal Inference (FCI) algorithm; (2) given the causal structure, we find a valid adjustment set and estimate the causal effects by an econometric model with the adjustment set as controlling non-causal effects. Our findings show that the spillover effects of WeChat are limited; in fact, only two other apps, Tencent News and Taobao, receive positive spillover effects from WeChat. In addition, we show that, if researchers fail to account for the causal structure that we determined from the graphical model, it is easy to fall into the trap of confounding bias and selection bias when estimating causal effects.

Chapter 2. ADOPTION OF ONLINE TWO-SIDED MARKET - TRANSPORTATION NETWORK COMPANY

2.1 INTRODUCTION

A transportation network company (TNC) is defined as “a service that does not own vehicles or employ drivers, and relies on software to connect passengers to rides” (California Public Utilities Commission, 2013). Early operators of transportation network apps include such well-funded firms as Uber, Lyft, Hailo, OlaCabs, and Didi Dache. A TNC creates a two-sided market with two versions of its app for passengers to generate orders and drivers to select orders. Both versions can be easily downloaded from major app platforms, including the App Store, Google Play, and Apps for Windows Phone.

A TNC app creates value in the following ways. For drivers, it provides enhanced functionality. For example, compared with traditional taxi drivers, TNC drivers can expand their customer pool beyond those who reserve by telephone or hail a taxi roadside and can select passengers based on the information about pickup location, destination, revealed tips, and other requirements to match their preferences for driving areas and routes. Another revolutionary imbedded function of the TNC app is that of allowing drivers to cancel orders, typically prohibited in the traditional reservation-based transportation model. With the canceling function, drivers can easily react to updated information about the quality of orders and about other options, such as roadside hailing. For passengers, a TNC app provides simplicity and flexibility in the transaction process. The most prominent function is “online-pay,” which enables passengers to handle their payments automatically. Once a trip is completed, the TNC app charges the passenger and immediately reports the fare, route, and time through email; then, if desired, the passenger is able to dispute the charge when he or she receives the email.

Another important role of a TNC app is to facilitate the launching of marketing campaigns. One typical challenge faced by TNC platforms is that potential users are reluctant to adopt this innovative online-to-offline-experience product due to their unfamiliarity with the features. Thus, companies use promotions. During the introduction of their apps, Uber and Lyft paid drivers a 5% bonus for every ride based on the app that they serviced. Recently, Uber raised one billion USD to promote the adoption of an Uber network in China, subsidizing new drivers as much as triple the regular fare per order. Other firms, such as Didi and Kuaidi in China, adopted similar strategies. A TNC app provides a channel through which the release, implementation, and adjustment of such a promotion can be conducted with ease. The availability of information of users' characteristics allows targeted promotion to specific segments of users, within well-controlled timeframes and locations, to complement the introduction of an app or even a specific function of the app. For example, a TNC can introduce a two-sided sales promotion for both drivers and passengers to receive cashback rewards for transactions conducted through online-pay to incentivize the use of this function. In addition to the direct effect of instantly increased usage, intensive promotion can generate, in the long run, the indirect effect of users' learning about the app, which may be perceived as a bundle of attributes, such as quality, price, and functionality. Through usage experience, users develop perceptions about the value of each attribute. A promotion can accelerate usage experience and help reduce user uncertainty about and increase recognition of the true value of each attribute, which, prior to use, is typically undervalued, according to the consumer learning literature.

To the best of our knowledge, there has been no study that systematically investigates and validates the effects of intensive promotion on two-sided online goods, such as TNC apps. Thus, in this paper, we aim to answer the following research questions: How does a two-sided sales promotion affect drivers' propensity to use a TNC app? How does a two-sided sales promotion interact with drivers' learning dynamically? How does a TNC design better promotion schemes to accelerate drivers' learning while being cost effective? We focus on the driver side of the TNC platform because the market in our context as well as in many others is in demand due to regulation; this makes the supply side more important to platform performance.

There are a number of challenges to measuring the effects of sales promotion on drivers. The first challenge arises from features such as order-canceling and online-pay. The remodeled business process, based on the use of the app, consists of sequential and interconnected decisions in each transaction, and sales promotion might have different effects on each transaction decision. The effect on one stage of a decision might indirectly affect another stage of the decision due to the link between them. For example, one sales promotion policy in our data sample is that both drivers and passengers receive a cashback bonus when passengers use the online-pay function. Therefore, an order with higher potential for a cashback bonus will increase the driver's willingness to accept and fulfill the order. Without understanding a multiple decision-making process enabled by new features, however, it is impossible to quantify the effects of sales promotion accurately. Therefore, we build a structural model to explicitly recover the data-generation process of drivers' sequential decisions in the presence of sales promotion.

In addition, because sales promotion during the introductory period might have indirect effects due to diminishing uncertainty, our model needs to account for uncertainties about the new features imbedded in different stages of decisions. A typical Bayesian learning model can handle this effectively; however, such a model is limited to the learning of one attribute associated with one decision. In our case, each usage experience consists of the attributes in multi-dimensional space, with each dimension's representing the attribute for a decision in one stage. Given that the decisions are related to each other, the learning associated with different decisions should be connected with each other. We construct our structural model with Bayesian learning conditional on earlier-stage decisions to model the connected learning processes of a driver's willingness to accept an app-generated order and to fulfill the order and the driver's belief about the passenger's willingness to use the online-pay function.

Using the data from a leading TNC in China, we quantify drivers' learning how to use the app to accept orders and that of use of the canceling function as well as drivers' beliefs about how passengers learn about the use of online-pay. We find that tips from passengers, subsidies from the app provider, and a cashback bonus for passengers and for drivers all affect drivers' decisions through not only the direct impact on the latent utility of taxi drivers but also drivers' beliefs about passengers' decisions to use online-pay. Our counterfactual analysis separates learning-induced indirect effects of sales promotion from the direct effects on usage increase. We also identify the effects of sales promotion for passengers on drivers' decisions and provide managerial suggestions for similar app providers on how to accelerate consumer learning while being cost effective.

The remainder of the paper is organized as follows. In Section 2.2, we briefly discuss the related literature and our contributions correspondingly. We describe the research context and our dataset in Section 2.3. In Section 2.4, we present our model of drivers' decisions to use the TNC app. In Section 2.5, we discuss our estimation strategy and report the estimation results. We simulate data to generate insights of our model and propose optimized sales promotion strategies in Section 2.6. Finally, we summarize our findings and conclude our research in Section 2.7.

2.2 LITERATURE REVIEW

Our research is related to the literature on the dynamics between empirical consumer learning and sales promotion. A large quantity of industry anecdotes identify sales promotion in the product introduction period as a strategy to enhance usage experiences and to foster consumer learning. In most research settings of experience goods with adequate variation of prices, however, sales promotion is limited in its direct effect and is considered equivalent to a temporary price cut. Consequently, very little attention has been paid to teasing out the indirect effects of sales promotion from the direct price effect. There are a few exceptions, for example, Erdem and Sun (2002), who find evidence of spillover effects of sales promotion and advertising in umbrella branding of multiple products. Chen et al. (2009) examine another extreme case of sales promotion as a permanent price cut for cigarettes and find the effectiveness of permanent price cut when considering forward-looking behavior, learning and addiction. In our empirical setting, we identify the indirect effects of sales promotion through consumer learning. In addition, we consider the format of two-sided promotion: an innovative marketing mix widely used in e-business. Our counterfactual analysis provides insight into how to improve such promotion.

With respect to methodological approaches, we apply a Bayesian learning model in a finite-horizon, forward-looking context. Erdem and Keane (1996) first identify consumer learning about product quality levels through experience and unobserved signals, such as advertising by applying the Bayesian updating process. Due to the learning model's applicability to consumers' choices under uncertainty, the model has been extended to account for many formats of information, such as learning from observed signals (Erdem et al. 2008), from online reviews with different credibility (Zhao et al. 2013) and with different weights, and from their own preference for multiple attributes and variance of preference (Wu et al. 2015). Researchers have applied the learning model in a more complex context, for example, pharmaceutical treatment (Crawford and Shum 2005, Chan and Hamilton 2006) and addictive products, such as cigarettes (Chen et al. 2009). In information systems research, the learning model has been used to understand content generation and consumption on the mobile Internet (Ghose and Han 2011), ideation on crowdsourcing (Huang et al. 2014), and online reviews (Ho et al. forthcoming). In our paper, we extend the Bayesian learning model by accounting for learning about multiple attributes associated with interrelated decisions. To our knowledge, this is the first information systems study that uses conditional Bayesian learning to address sequential decisions. To model sequential decisions, we follow Arcidiacono (2005), who investigates affirmative action's effects on sequential decisions in a higher education application.

Our paper is also the first to investigate the micro level of on-demand TNC apps from the information systems perspective. Even though TNC apps have attracted a great deal of attention in industry and the media, there is limited research on this topic. Some preliminary work focuses on the impact of this newly introduced product on the traditional taxi industry and public transportation systems. Rayle et al. (2014) find that ride-sourcing complements traditional taxis and public transit by introducing younger passengers, when competing with traditional taxis and public transit. They also find that ride-sourcing results in a shorter wait time for traditional taxis and in a shorter overall delivery time in regard to public transit. By constructing a model of costs of transaction and regulation, Li et al. (2014) show that total cost and stakeholders' cost decrease by utilizing taxi apps and suggest the use of taxi apps to replace current the taxi-calling system, with centralized management from government. In a study that departs from research on traditionally assessed efficiency benefits, Li and Zhao (2015) conducted interviews with drivers, passengers, regulators, association leaders, app providers, and developers and conclude that TNC apps can reduce the taxi usage-seeking behavior of dispatchers, while humanizing the relationship between drivers and passengers. Recently, in a case study of a taxi app, Tan et al. (2015) find, through a situational and artefactual affordances approach, that a taxi app can use gamification as the implementation of digital technology-enabled changes to better meet customer needs. None of the prior literature, however, has explained the TNC app business model from a micro-level perspective, as we do in this study. We model the transactional level behavior of TNC app users and empirically identify hidden utility-level parameters through a structural approach.

2.3 Research Context and Data

Our data are from a TNC that is the leading mobile app-based transportation network in China, with over 150 million registered passengers, 1.5 million registered taxi drivers, and 5 million transactions per day. The structure of their app offers an ideal means for studying the value of the app's attributes. Different from the Uber app, which generates additional transportation, this app applies only to the existing taxis on the market due to government regulation of this market. In other words, there are no new taxis introduced by the app, and all app users on the supply side are traditional taxi drivers. This helps to control for other effects, such as car characteristics and driver characteristics, leaving the introduction of the app and promotional activities as the only exogenous changes. Each driver makes a decision on whether to use an app or a traditional taxi business model when their taxis are empty.

In addition, the sample that we analyzed is from a city where the supply of taxis is significantly lower than the demand. The statistics show a ratio of more than 500 people per taxi in this city. These statistics are consistent with the popularity of the app; based on the usage pattern in our sample, orders arrive every few minutes. This erases the concern about endogeneity due to omitted demand when we model drivers' decisions to use the app or to take orders in a traditional fashion.

From a driver's perspective, each transaction involves three stages. Prior to Stage 1, the app provider announces cashback plans to drivers and passengers. Note that the policy is to give cash back to drivers and passengers only if passengers pay the fare online instead of by cash. Similar to most marketing campaigns, the cashback plan for passengers is made public but that for drivers is revealed to drivers only through the app. In other words, the drivers' information set consists of cashback plans for drivers and passengers, whereas the passengers' set is limited to the plan for themselves only. Given the policy, an order is initiated by a passenger who needs a taxi and who inputs information on his or her current location and destination. Usually, this step is automated by GPS and voice messages, unobservable to econometricians. The passenger could additionally reveal tips to the driver to increase the probability that his or her order will be accepted. After that, the platform also reveals the subsidy for the driver, which is at a constant rate for several short-term rides, regardless of whether the passenger pays online. Both tips and subsidy are revealed to the driver before he or she accepts the order. This is different from the traditional business model. Subsidies and tips are paid independently and added to cash back and the regular fare. The order request will be sent, upon the completion of the inputs, to all available and nearby drivers.

In Stage 1, a driver will decide, given the information noted above, whether to accept an order from the app. If a driver accepts an order from the app, he or she will receive the regular fare, the passenger's tip, and a platform subsidy (if applicable); if the passenger pays online, the driver also will receive a cashback bonus. If the driver does not accept the order, then he or she will take orders from other sources, e.g., roadside taxi hailers, the call center, or wait for the next order from any source. In this regard, it should be noted that app-based orders are not necessarily superior to regular orders. Even though drivers can receive benefits, such as monetary rewards as well as ease of use due to the app's functionality, they also incur extra cost for the app's functions. For example, drivers need to use third-party financial services to process transactions, which reduces the instant gratification associated with cash. There are also some additional costs related to the use of the app, such as drivers' efforts to install the app and to pay attention to notifications from the app as well as the idle time while driving to pick up passengers.

After accepting an order, in Stage 2, a driver can decide whether to fulfill or cancel it. A driver cancels if he or she deems the order as low quality or risky or if better alternatives are available. The decision to cancel an order is similar to that of accepting an order and depends on time or location. A driver may receive updated information on an order that he or she accepted that may trigger a change of mind. In this way, a low-quality or risky order that may have been abandoned in the traditional business model might be accepted first and then canceled later, given the updated information. The availability of this choice will result in a distinctive learning pattern for drivers.

Stage 3 involves the decision by a passenger to pay online, conditional on the ride's being completed. Paying online incurs a positive attribute, such as convenience, as well as potential negative ones; for example, the function may be difficult to use, is unstable, or has a high failure rate. To promote the use of online-pay, in our context, cashback awards for both sides are contingent on the decision of a passenger to pay online. Even though this is not his or her own decision, a rational driver will form a belief about this decision because he or she may benefit from it.

Rational drivers are both backward-looking and forward-looking when using a TNC app. Backward-looking means that a driver learns to use the different functions associated with the different stages, based on his or her usage experience. Each usage experience signals to the driver the attributes of functions such that he or she will be more certain when using the app. Forward-looking implies that a driver makes decisions in earlier stages based on his or her belief about potential outcomes in later stages as a means to optimize his or her overall utility. This behavior connects the decisions made in different stages.

We collect a panel of structured data from the TNC firm. Our data include 952 taxi drivers who have used the TNC app at least once during the introductory period. We keep track of all the transactions of this set of drivers since their registration. In total, we have 198,689 transactional records. We observe the following seven variables: whether drivers accept the orders, whether transactions are fulfilled, whether online pay is used, cashback amount for passengers when using online-pay, cashback amount for drivers when passengers use online-pay, tips from passengers, and subsidies from the TNC platform. These data, in general, provide a record of the business process and sales promotion by the TNC platform. We show summary statistics in Table 2.1.

Table 2.1. Summary Statistics

Variable	Mean	SD	Min	Max
Acceptance	0.7622	0.4963	0	1
Fulfillment	0.6388	0.4803	0	1
Online-Pay	0.4394	0.4257	0	1
Cashback for Driver*	1.4733	0.3625	0	1.5625
Cashback for PSGR*	1.5843	0.6768	0	2.5000
PSGR Tips*	0.0076	0.1531	0	7.8125
Platform Subsidies*	0.0300	0.2149	0	15.6250

*The unit of currency is the US dollar

To understand the effects of sales promotion and of learning, we analyze daily aggregate level transaction amounts that correspond to daily promotion policies, as Figure 2.1 shows. The time horizon in the figure covers three months in the introductory period of the app. The daily average outcome variables are shown with lines in black, red, and blue to represent daily acceptance rate, daily fulfillment rate, and daily online-pay rate, respectively. In general, usage of online-pay and fulfilled transactions are proportional to the accepted amounts, with a rising trend over time. The sales promotion policy chart shows that cashback promotion policy also changed over time. For drivers, it started at 0 and, by Day 10, had increased, an increase that, with some fluctuation, was maintained. For passengers, it also started at 0, went up from Day 10 to Day 70, and declined afterward. The other rewards chart shows a declining trend of tips from passengers and subsidies from the platform.

Figure 2.1. Sales Promotion vs. Outcome Variable Dynamics

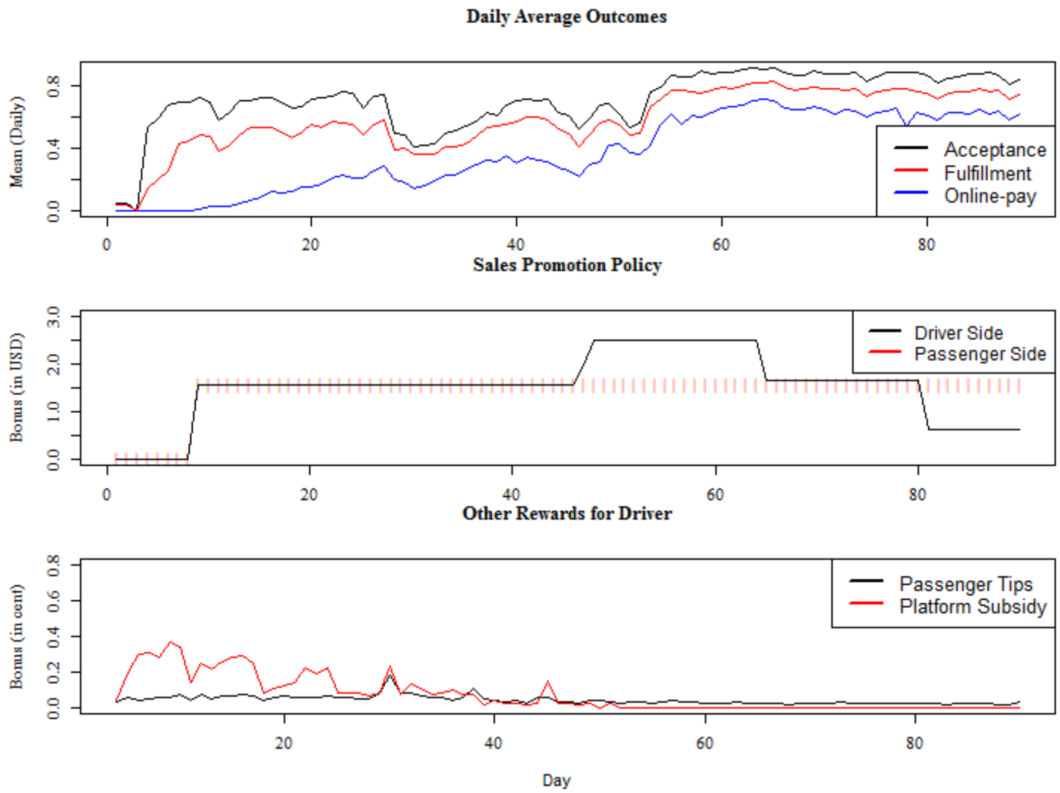

Table 2.2. Change Over Time

Statistics	Outcomes				Sales Promotion and Other Rewards			
	#Orders	Accept	Fulfill | Accept	Online-pay | Fulfill	Cash Back for PSGR*	Cash Back for Driver*	Platform Subsidy*	PSGR Tips*
Overall	198689	0.76	0.84	0.69	1.58	1.47	0.04	0.03
First 30	56291	0.62	0.71	0.28	1.25	1.25	0.19	0.06
Mid 30	67259	0.77	0.87	0.74	2.20	1.56	0.01	0.03
Last 32	75139	0.86	0.88	0.82	1.28	1.56	0	0.02

*The unit of currency is the U.S. dollar

When comparing the outcome variables with respect to sales promotion policy, we note several interesting observations. We find overall increases with more intensive sales promotion, which indicates that sales promotion stimulates acceptance, fulfillment, and online-pay rates. In addition, we find that all of the outcomes, including acceptance and fulfillment rates, which are controlled by drivers, increase when the cashback for passengers is raised, around Day 50. This implies that sales promotions for passengers might affect driver decisions. Further, a comparison of the outcomes after Day 70 with those before Day 45 shows that, even though sales promotion policies for both sides are almost the same, and tips from passengers and subsidies from the platform are even higher during earlier period, more orders are accepted and fulfilled with online-pay after Day 70. One potential explanation for the difference between those two periods is that the usage experience accumulated between Days 45 and 70 leads drivers to fully perceive the value of the app and converts them to frequent users. A similar pattern can be observed by dividing the time horizon into three even periods and comparing the average outcomes, as shown in Table 2.2. Comparing the middle 30 days with the last 32 days, all three outcome variables are improved, even with declining sales promotions and other rewards.

2.4 MODEL

We present a structural model of the data-generation process of how a typical order on a TNC app is accepted, fulfilled, and finished following a request from a passenger. As discussed earlier, an order will go through three stages of decisions:

Stage 1: The driver decides whether to accept the order;

Stage 2: The driver decides whether to cancel the accepted order;

Stage 3: The passenger decides whether to redeem the sales promotion by paying online.

Note that the driver's decision in Stage 1 is conditional on his or her belief in the decisions in Stages 2 and 3 and that the decision in Stage 2 is conditional on that in Stage 3. That is, the model needs to account for finite-horizon, forward-looking dynamics that incorporate the relationship among different stages with a decision-dependent state transition. We model these decisions in a forward way for easier interpretation but estimate the model by following typical backward induction. Specifically, we subscript the parameters associated with the decision to accept an order with a, of fulfilling an accepted order with s, and of redeeming sales promotion with c, to avoid potential confusion. We further define the arrival of each order as a one-time period.

2.4.1 *Stage 1: Decision to Accept an Order*

When an order arrives, the driver makes the decision about whether to use the TNC app to accept or to decline and chooses an alternative from traditional sources, e.g., telephone call, airport pickup, hotel pickup, passengers on the roadside. We assume that driver i at time t will receive a utility of u_{a1t} if he accepts an order from the TNC app, and u_{a0t} if not. The subscript for individuals, i, is suppressed, given the homogeneous assumption of taxi drivers who are all employed by taxi companies.

Two possible outcomes will occur after a driver accepts an order from the app. If the driver is committed to the order, the transaction will be fulfilled, and he or she will gain utility, including fare, tips, subsidy, and sales promotion, less the cost. In contrast, if the driver deems an order unprofitable or risky, he or she will cancel the order. By using an indicator function, we model the utility of accepting an order as:

$$u_{a1t} = D_{st}U_{s1t} + \left(1 - D_{st}\right)U_{s0t} + \varepsilon_{a1t}, \tag{2.1}$$

where D_{st} is a dummy variable that indicates whether the transaction is fulfilled, U_{s1t} is the utility if the order is fulfilled, and U_{s0t} is that if the order is canceled. We specify ε_{a1t} to capture unobserved information by econometricians, such as location, which is assumed to follow a type-I extreme value distribution.

Other than unobserved information, there are three components that account for the driver's utility when the transaction is fulfilled, as specified in (2). The driver gains utility from the attribute of using the app to accept an order and all forms of bonus, net the cost, if online-pay is used for that order. Here, the Stage 1 specific attribute A_a refers to an aggregation of product-specific characteristics that would generate utility to drivers when they choose to use the app to accept orders. The second component includes the three formats of bonus for taxi drivers. The first component is cashback bonus B_t^{cb} if passengers use online-pay, which is used for the purpose of promoting use of the app and online-pay function, redeemed from the TNC platform and its financial partners. The second one is passenger tips B_t^{tip}, which is revealed before drivers make a decision, to increase drivers' willingness to accept specific orders. The third is platform subsidy B_t^{sub}, which is similar to passenger tips and is used to encourage drivers to accept orders but is provided by the TNC platform. Despite these monetary rewards, using online-pay may incur potential transaction costs as well as the loss of instant gratification. We use the additive form with scalars for attribute, monetary unit of bonus, and online-pay associated cost as follows:

$$U_{s1t} = \beta_{a1} A_a + \beta_{a2} \left(D_{pt} \times B_t^{cb} + B_t^{tip} + B_t^{sub} \right) + D_{pt} c_{a1}, \tag{2.2}$$

where β_{a1} represents weight of utility from attribute A_a of the app. We specify a constant weight β_{a2} to monetary rewards in different forms because they are the same in units as well as in formats.

D_{pt} is a dummy variable that indicates whether the passenger uses online-pay, and c_{a1} is the cost associated with online-pay.

When a transaction is canceled, the driver still receives attribute value A_a associated with use of the app to accept the order, as the usage associated with accepting an order already has occurred. In addition, he or she receives the utility associated with canceling the order c_{a2}, as the cancelation prevents the further cost of a risky order. We formulate the utility of canceling an order as follows:

$$U_{s0t} = \beta_{a1} A_a + c_{a2}. \tag{2.3}$$

Taken together, we have the utility of accepting an order as:

$$u_{a1t} = \beta_{a1} A_a + D_{st} \left(\beta_{a2} \left(D_{pt} B_t^{cb} + B_t^{tip} + B_t^{sub} \right) - D_{pt} c_{a1} \right) + \left(1 - D_{st}\right) c_{a2} + \varepsilon_{a1t}. \tag{2.4}$$

A driver could decide not to accept an order when he or she receives a higher utility from outside goods. We model the utility of outside goods u_{a0t} to be the summation of constant level c_{a0} plus a stochastic error term that captures unobserved utilities:

$$u_{a0t} = c_{a0} + \varepsilon_{a0t}. \tag{2.5}$$

Given that the TNC platform has been recently introduced to the market, drivers are not initially fully certain about the latent utility. Therefore, before a driver makes the decision of whether to use the app to receive orders, he or she first forms an expectation of the utility from two alternatives, based on the information updated to period t, and chooses the alternative that maximizes his or her expected utility. The major source of information that a driver uses to improve his or her perception of utility is his or her own usage experience, which provides more precise information than do other channels. We define information set I_t as the cumulative usage experience. By taking expectation conditional on I_t and rearranging our expected utility function by linearity of conditional expectation, we have the following expressions for conditional expectation of latent utility for the two alternatives:

$$E\left(u_{a1t} \mid I_t\right) = \beta_{a1} E\left(A_a \mid I_t\right) + \Pr\left(D_{st} \mid I_t\right)\left(\beta_{a2}\left(\Pr\left(D_{pt} \mid D_{st}, I_t\right) B_t^{cb} + B_t^{tip} + B_t^{sub}\right) - \Pr\left(D_{pt} \mid D_{st}, I_t\right) c_{a1}\right)$$
$$+ \left(1 - \Pr\left(D_{st} \mid I_t\right)\right) c_{a2} + \varepsilon_{a1t}, \tag{2.6}$$

$$E\left(u_{a0t} \mid I_t\right) = c_{a0} + \varepsilon_{a0t}. \tag{2.7}$$

In modeling this way, we let the uncertainty be absorbed into three components: the expected value of attribute $E\left(A_a \mid I_t\right)$, the probability of fulfilling an order $\Pr\left(D_{suc} \middle| B_{it}^{bonus}, C_{it}^{bonus}, B_{it}^{reward}, B_{it}^{subsidy}, I_{it}\right)$, and the probability that a passenger pays online $a_2(\Pr\left(D_{onlinepay} \middle| B_{it}^{bonus}, C_{it}^{bonus}, B_{it}^{reward}, B_{it}^{subsidy}, D_{suc}, I_{it}\right)$. We assume that drivers behave as Bayesian learners, who update their expectations based on I_t. Specifically, because drivers might get involved in three sequential decisions and receive usage experience from each of them, we let drivers update specific decisions based on the usage experience of a corresponding decision before time t. Therefore, $\Pr\left(D_{st} \mid I_t\right)$ and $\Pr\left(D_{pt} \mid D_{st}, I_t\right)$ are updated with Stage 2 and Stage 3 decisions, respectively, which will be explained in later sections. $E\left(A_a \mid I_t\right)$ is updated with the Stage 1 decision that drivers experience from the beginning to time t.

We explicitly explain how $E\left(A_a \mid I_t\right)$ is formed through a Bayesian learning process. Before starting to use the app to accept orders, drivers have prior information about the attribute of accepting orders from the app. We model the prior information following $N\left(A_{a0}, \sigma_{a0}^2\right)$ to accommodate a potentially biased prior belief A_{a0} and$\mathrm{N}(\mathrm{A}_0, \sigma_{\mathrm{A}0}^2)$ drivers' uncertainty σ_{a0}^2. Drivers make first-time decisions based on prior information only, such that first-time attribute A_a is drawn from prior distribution $N\left(A_{a0},\sigma_{a0}^{2}\right)$. By defining $\sigma_{A_a 0}^{2}$ as the variance of driver i's perception of the mean attribute level at the very beginning, we have:

$$E\left(A_a \mid I_0\right)=A_{a0}, \text{ and } \sigma_{A_a 0}^{2}=Var\left(A_a \mid I_0\right)=E\left(\left(A_a-A_{a0}\right)^{2} \mid I_0\right)=\sigma_{a0}^{2}. \tag{2.8}$$

Drivers update their beliefs about attribute value and uncertainty when they receive signals, and such signals, with some noise, are drivers' own usage experience by which they can perceive the true attribute value of using the app to accept orders. The noise can be derived from the variability of true attribute value itself or the variability associated with specific context in usage experience. To make the Bayesian update conjugate, we assume that the signal of the app's true attribute value, denoted as A_t^e, follows a normal distribution, according to:

$$A_t^e \sim N\left(A_{a1},\sigma_{a1}^{2}\right). \tag{2.9}$$

where A_{a1} $\mathrm{N}(\mathrm{A}_1,\sigma_{\mathrm{A}}^2)$. A_1 is the mean of the signal that equals to true attribute value, and σ_{a1}^{2} captures the variance of the signal.

We model that a driver updates if he or she receives one more usage experience or, if not, stays with the initial perception. Specifically, when a driver experiences use of the app to take an order at time $t-1$, he or she updates his or her perception as a weighted average of the perception formed in the last time period $E\left(A_a \mid I_{(t-1)}\right)$ and the newly received signal A_{t-1}^e. To be consistent with the intuition that a more precise signal leads a driver's perception to be closer to the true attribute value, we model the weights as precision parameters, using the inverse of perception variance and that of the signal, according to:

$$E\left(A_a \mid I_t\right)=D_{at}\times\frac{\dfrac{A_{t-1}^e}{\sigma_{a1}^{2}}+\dfrac{E\left(A_a \mid I_{(t-1)}\right)}{\sigma_{A_a(t-1)}^{2}}}{\dfrac{1}{\sigma_{a1}^{2}}+\dfrac{1}{\sigma_{A_a(t-1)}^{2}}}+\left(1-D_{at}\right)\times E\left(A_a \mid I_{(t-1)}\right), \tag{2.10}$$

$$\sigma_{A_a t}^2 = D_{at} \times \frac{1}{\frac{1}{\sigma_{a1}^2} + \frac{1}{\sigma_{A_a(t-1)}^2}} + \left(1 - D_{at}\right) \times \sigma_{A_a(t-1)}^2, \tag{2.11}$$

where D_{at} is a dummy variable that indicates whether the driver accepts an app order. Posterior uncertainty is updated as the inverse of the sum of the inverse of prior uncertainty and the inverse of signal variance if the signal is received. The Bayesian updating rule above exhibits diminishing uncertainty with gaining usage experience, and faster uncertainty diminishing with gaining less noisy signals in usage experience.

Rational drivers form beliefs about $\Pr\left(D_{st} \mid I_t\right)$ and $\Pr\left(D_{pt} \mid D_{st}, I_t\right)$, following the rules in Stage 2 and Stage 3, conditional on the corresponding usage experience until time t. If we assume that $\Pr\left(D_{st} \mid I_t\right)$ and $\Pr\left(D_{pt} \mid D_{st}, I_t\right)$ are formed and that the error terms ε_{a0t} and ε_{a1t} are independently and identically distributed with type-I extreme distribution, we obtain the probability that a driver accepts an order from the TNC app conditional on information I_t and the stage-specific log likelihood function as follows:

$$\Pr\left(D_{at} \mid I_t\right) = \frac{\exp\left(U_{at}\right)}{\exp\left(U_{at}\right) + \exp\left(c_{a0}\right)}, \tag{2.12}$$

where

$$\begin{aligned} U_{at} = {} & \beta_{a1} E\left(A_a \mid I_t\right) + \Pr\left(D_{st} \mid I_t\right)\left(\beta_{a2}\left(\Pr\left(D_{pt} \mid D_{st}, I_t\right) B_t^{cb} + B_t^{tip} + B_t^{sub}\right) - \Pr\left(D_{pt} \mid D_{st}, I_t\right) c_{a1}\right) \\ & + \left(1 - \Pr(D_{st} \mid I_t)\right) c_{a2}, \end{aligned} \tag{2.13}$$

$$L_1(\beta_a, \beta_s, \beta_p) = \sum_{i=1}^{n} \sum_{t=1}^{T_i} \left(\log \Pr\left(D_{at} \mid I_t\right) D_{at} + \log\left(1 - \Pr\left(D_{at} \mid I_t\right)\right)\left(1 - D_{at}\right)\right). \tag{2.14}$$

2.4.2 *Stage 2: Decision to Fulfill or Cancel an Order*

After a driver has accepted an order from the TNC platform, he or she still has the opportunity to cancel if he or she obtains more information that signals low quality of the order or if better alternatives appear. We assume driver i at time t will receive utility of u_{s1t} if he or she fulfills an order from the TNC app and u_{s0t} if not.

There are six systematic components that constitute a driver's utility when he or she decides whether to fulfill an order. These components include a cashback bonus B_t^{cb} if the passenger uses online-pay, passenger tips B_t^{tip}, platform driver subsidy B_t^{sub}, cost associated with online-pay c_{s1}, an initially uncertain attribute value of the canceling feature A_s if a driver is to fulfill the order, and the utility for outside goods modeled as a constant c_{s0}. We introduce error terms ε_{s1t} and ε_{s0t} , respectively, for fulfilling an order and canceling an order to represent additional unobserved information. We specify latent utility of fulfilling or canceling an order as a linear additive form:

$$u_{s1t} = \beta_{s1}A_s + \beta_{s2}B_t^{sub} + \beta_{s3}B_t^{tip} + D_{pt}\left(\beta_{s4}B_t^{cb} + c_{s1}\right) + \varepsilon_{s1t}, \tag{2.15}$$

$$u_{s0t} = c_{s0} + \varepsilon_{s0t}, \tag{2.16}$$

where D_{pt} is an indicator function that governs the utility associated with using online-pay. Only when online-pay is used can the driver gain cash back from the platform, while incurring a transaction cost. β_{s1} is the weight for an attribute value; β_{s2} captures the weight for a platform subsidy; β_{s3} is the weight for a revealed tip from a passenger; and β_{s4} represents the weight for online-pay contingent cash back. We give different weights to different formats of monetary rewards because we suspect that a revealed tip, as well as a subsidy, might play an additional role of signaling the quality of an order. We conjecture that a high tip might convey extra information from a passenger, for example, potential extra cost because the tip might indicate that solely flat rate netting the cost is not as competitive as those in alternative orders for drivers. Similarly, an order-specific subsidy from the platform also might convey the information known by the platform, whereas such an effect does not exist for cashback bonus. It is not modeled in the same way in Stage 1, as the same information already is captured by $\Pr(D_{st} \mid I_t)$, which avoids the overlapped information and helps to identify the cost to cancel c_{a2}.

Because the canceling function has been recently introduced, drivers are uncertain about how this function works and are concerned that they might overact or underact in regard to canceling orders. To rationalize our model, we assume that a driver will form the expectation of latent utility conditional on the information up to time t to make the decision to cancel an order. Given that monetary rewards are revealed precisely before the decision, all the uncertainties in latent utility stem from A_s, as an aggregation of characteristics of the canceling feature that is most likely to be uncertain for drivers, as well as D_{pt}, which captures the uncertain level associated with next-stage decision. Accordingly, the expected utility conditional on information set I_t is given by:

$$E\left(u_{s1t} \mid I_t\right) = \beta_{s1} E\left(A_s \mid I_t\right) + \beta_{s2} B_t^{sub} + \beta_{s3} B_t^{tip} + \Pr\left(D_{pt} \mid D_{st}, I_t\right)\left(\beta_{s4} B_t^{cb} + c_{s1}\right) + \varepsilon_{s1t}, \quad (2.17)$$

$$E\left(u_{s0t} \mid I_t\right) = c_{s0} + \varepsilon_{s0t}. \quad (2.18)$$

A driver needs to form beliefs about the attributes of fulfilling an order through his or her experience. Similar to the model of the perceived attribute value of accepting orders from the app, we model these learning processes by following the Bayesian updating rule with the prior perceived value as following $\mathrm{N}(\mathrm{S}_0, \sigma_\mathrm{S}^2)$ and the signal as following $N\left(A_{s1}, \sigma_{s1}^2\right)$. We suppress the updating rules here for a concise interpretation and let$\mathrm{E}(\mathrm{S}_{\mathrm{waitit}} | \mathrm{C}_{\mathrm{it}}^{\mathrm{bonus}}, \mathrm{B}_{\mathrm{it}}^{\mathrm{bonus}}, \mathrm{I}_{\mathrm{it}})$. represent the Bayesian-updated driver's belief of perceived attribute value for fulfilling an order at time t. In addition, drivers form beliefs about passengers' willingness to use online-pay as $\Pr\left(D_{pt} \mid D_{st}, I_t\right)$, conditional on usage experience until time t by following the rule seen in Stage 3. We simply model the error term following type-I extreme value distribution, which results in a closed-form logit formula for the probability of fulfilling an order, conditional on the information set. We present the probability of fulfilling an order as well as the Stage 2 log likelihood function as follows:

$$\Pr(D_{st} \mid I_t) = \frac{\exp\left(\beta_{s1} E\left(A_s \mid I_t\right) + \beta_{s2} B_t^{sub} + \beta_{s3} B_t^{tip} + \Pr\left(D_{pt} \mid D_{st}, I_t\right)\left(\beta_{s4} B_t^{cb} + c_{s1}\right)\right)}{\exp\left(\beta_{s1} E\left(A_s \mid I_t\right) + \beta_{s2} B_t^{sub} + \beta_{s3} B_t^{tip} + \Pr\left(D_{pt} \mid D_{st}, I_t\right)\left(\beta_{s4} B_t^{cb} + c_{s1}\right)\right) + \exp\left(c_{s0}\right)}, \quad (2.19)$$

$$L_2(\beta_s, \beta_p) = \sum_{i=1}^{n} \sum_{t=1}^{T_{ai}} \left(\log \Pr\left(D_{st} \mid I_t\right) D_{st} + \log\left(1 - \Pr\left(D_{st} \mid I_t\right)\right)\left(1 - D_{st}\right)\right). \quad (2.20)$$

2.4.3 *Stage 3: Passenger Decision to use Online-Pay and Redeem a Sales Promotion*

Conditional on an order's being fulfilled, the passenger makes the decision of whether to use online-pay for taxi fare. Given that the online-pay function was recently introduced to passengers, the entire population of passengers can be considered Bayesian learners. Similarly, through learning that is accelerated by the cashback policy, the perceived attributes associated with the online-pay function evolves dynamically and eventually converges to a stationary level from a biased starting value.

Given that online-pay could potentially lead to a cashback bonus from the platform, a driver calculates his or her expected utilities in Stages 1 and 2 with the expectation of the order's being paid online. Different from learning from his or her own experience of using a function of the app in Stages 1 and 2, in Stage 3, a driver forms his or her belief about receiving an online-paid order through interactions with his passengers. In other words, a driver's perception of the probability of online-pay reflects the aggregate usage pattern of all his passengers. Hence, from the perspective of structural modeling, the probability of online-pay serves as a state variable for the driver, and the state transition (or the probability change) is dictated by passengers' decisions to redeem a cashback bonus.

To form his or her belief about the probability of online-pay, a rational driver formulates a representative passenger's decision problem. It is known to a driver that the passenger's information set is limited to passenger-side sales promotion C_t^{cb}, a tip promised when initiating the order B_t^{tip}, the attribute value of using online-pay function A_p, and the utility from outside alternative payment methods c_{p0}, such as cash or a prepaid Metro cards. These components, except the utility of outside goods, constitute a representative passenger's latent utility function of using online-pay. It follows a linear additive form, with β_{p1}, β_{p2}, and β_{p3} as the corresponding weights. To capture unobserved information, type-I extreme distributed error terms are included. Accordingly, the utility functions of taking online-pay and alternative way are given by:

$$u_{p1t} = \beta_{p1} A_p + \beta_{p2} C_t^{cb} + \beta_{p3} B_t^{tip} + \varepsilon_{p1t}, \tag{2.21}$$

$$u_{p0t} = c_{p0} + \varepsilon_{p0t}. \tag{2.22}$$

Passengers will gain cash back if they use the online-pay feature, and we conjecture that cash back for passengers incentivizes using online-pay. Revealed tips, in contrast, incur a cost for passengers, regardless of paying online or offline. Similarly, the expected utility and, hence, a driver's belief in the probability of receiving an online-paid order conditional on time information set I_t are:

$$E\left(u_{p1t} \mid D_{st}, I_t\right) = \beta_{p1} E\left(A_p \mid I_t\right) + \beta_{p2} C_t^{cb} + \beta_{p3} B_t^{tip} + \varepsilon_{p1t}, \tag{2.23}$$

$$E\left(u_{p0t} \mid D_{st}, I_t\right) = c_{p0} + \varepsilon_{p0t}, \tag{2.24}$$

$$\Pr\left(D_{pt} \mid D_{st}, I_{it}\right) = \frac{\exp\left(\beta_{p1} E\left(A_p \mid I_t\right) + \beta_{p2} C_t^{cb} + \beta_{p3} B_t^{tip}\right)}{\exp\left(\beta_{p1} E\left(A_p \mid I_t\right) + \beta_{p2} C_t^{cb} + \beta_{p3} B_t^{tip}\right) + \exp\left(c_{p0}\right)}. \tag{2.25}$$

The perceived aggregate attribute of using online-pay, A_p, is uncertain and needs to be learned by drivers. We model that drivers' prior belief of the attribute follows $N\left(A_{p0}, \sigma_{p0}^2\right)$ and the received signal about the attribute follows $N\left(A_{p1},\sigma_{p1}^{2}\right)$. Every time a driver fulfills an order, and the passenger uses online-pay for the fare, the driver's perception of this attribute is updated following the Bayesian rule, similar to what is described for Stage 1. Following the literature on structural models by Rust (1987), Hotz and Miller (1993), and Bajari et al. (2007), we assume that drivers are able to form consistent beliefs about the population-level behavior of using online-pay and that the state transition due to passengers' decision to use online-pay can be consistently inferred by drivers. This assumption allows us to link realized decisions of passengers with perceived probability of drivers to identify the associated parameters as follows:

$$L_3(\beta_p)=\sum_{i=1}^{n}\sum_{t=1}^{T_{si}}\left(\log\Pr\left(D_{pt}\mid D_{st},I_t\right)D_{pt}+\log\left(1-\Pr\left(D_{pt}\mid D_{st},I_t\right)\right)\left(1-D_{pt}\right)\right). \tag{2.26}$$

2.5 Estimation Results

2.5.1 *Identification*

We briefly discuss the identification of our model in two steps. In the first step, we explain the identification of a typical Bayesian learning process, following Crawford and Shum (2005). In the second step, we explain the identification of our full model.

There are three sets of parameters in a typical Bayesian learning model for discrete choice, with only one uncertain attribute value. The first and most intuitive one is the coefficients associated with exogenous variables, such as coefficients for different forms of sales promotion in our study. These coefficients are identified by the variation of variables associated with them. The second set of variables includes true attribute value, prior attribute value, and constant terms in each Bayesian updating process, among which only two can be identified, as only the difference matters in a discrete-choice model. Our learning pattern allows identification of the difference between prior and true value from the difference of decisions between earlier and later stages. In other words, the variation of different decisions and variation along time allow us to identify two parameters. Because we are more interested in the difference between true and prior attribute value and whether the true attribute value is higher or lower than the prior attribute value, we normalize the true attribute value to zero and leave the prior attribute value and constant with freedom. This is also the most computationally effective way for minimization algorithm.

The third set of parameters are prior variance and true variance of attributes for each Bayesian learning process, such as σ_{a0}^2 and σ_{a1}^2, σ_{s0}^2 and σ_{s1}^2, and σ_{p0}^2 and σ_{p1}^2. We are allowed to identify only one of prior variance and true variance in each Bayesian updating because only the relative difference between those two variance parameters matters. Given a fixed prior mean value as the starting point and a fixed true mean value as the converged endpoint, the latent utility still has freedom in the speed of convergence as well as with the time dimension. This convergence rate can be visualized as a curvature along the time series and identifies relative variance between prior variance and posterior variance. In our model, we normalize true value variance parameters as 10 and identify the prior variance parameters to investigate the uncertainty levels of attributes before any usage experience.

The identification for the full model is straightforward, following our discussion above about the identification for Bayesian learning. Given the observations of three sequential decisions as drivers' decisions to accept orders, drivers' decisions to fulfill orders, and passengers' decisions to pay online in drivers' perception, we can identify Bayesian updating associated parameters A_{a0}, σ_{a1}^2, A_{s0}, σ_{s1}^2, A_{p0}, and σ_{p1}^2 from the difference of latent perceived attributes between earlier time points and later time points and the curvature of perceived attributes for each of the decisions.

Other parameters are components in a generalized linear formula, which follows typical identification rules of discrete-choice models. Given the data of three sequential decision choices, we are able to identify linear components by the difference in latent utility between different choices when a type-I extreme value error is used, which imposes fixed variability of error terms. One concern could be the identification of parameters associated with the probability of online-pay and of canceling in the Stage 1 decision, as they are functions of sales promotion, and some sales promotion variables appear in different decisions. This concern can be erased by Bayesian learning in later stages, as the mechanism of Bayesian learning itself introduces additional information from past experience with decisions. With these exclusive variables and logit transformation, which imposes a non-linear transformation of variables, our model achieves full identification.

2.5.2 *Estimation Specification and Model Fit*

We use a simulated maximum likelihood estimation method to recover parameters by following a two-stage estimation method. With independent errors across the stages, the log likelihood function can be divided into three terms:

$$L\left(\beta_a,\beta_s,\beta_p\right)=L_3\left(\beta_p\right)+L_2\left(\beta_s,\beta_p\right)+L_1\left(\beta_a,\beta_s,\beta_p\right), \tag{2.27}$$

$L_3\left(\beta_p\right)$ - the log likelihood contribution of paying online (redeem cash back);

$L_2\left(\beta_s,\beta_p\right)$ - the log likelihood contribution of fulfilling an order;

$L_1\left(\beta_a,\beta_s,\beta_p\right)$ - the log likelihood contribution of accepting an order;

$\beta_p=\left\{A_{p0},\log\left(\sigma^2_{p1}\right),\log\left(\beta_{p1}\right),\beta_{p2},\beta_{p3},c_{p0}\right\}$

$\beta_s=\left\{A_{s0},\log\left(\sigma^2_{s1}\right),\log\left(\beta_{s1}\right),\beta_{s2},\beta_{s3},\beta_{s4},c_{s1},c_{s0}\right\}$

$\beta_a=\left\{A_{a0},\log\left(\sigma^2_{a1}\right),\log\left(\beta_{a1}\right),\beta_{a2},c_{a2},c_{a1},c_{a0}\right\}$

Here β_p are state parameters, and β_s, β_a are utility parameters. Consistent estimates of β_p can be identified from maximizing $L_3\left(\beta_p\right)$. With the estimates of β_p, consistent estimates of β_s can be found by optimizing $L_2\left(\beta_s, \beta_p\right)$. Finally, coupled with estimates of β_s, β_p, estimates of β_a can be obtained by maximizing $L_1\left(\beta_a, \beta_s, \beta_p\right)$.

The likelihood maximization in Stage 3 now reduces to a typical Bayesian learning model. More specifically, we use the typical maximal simulated likelihood method proposed by Erdem and Keane (1996) to recover state transition parameters in Stage 3, with the attribute perception as simulated and the variance as integrated numerically. Because state transition is decision specific, and, in our case, if an order is canceled, there will be no online-pay or corresponding learning, so we need only the data in which the Stage 2 decision is not to cancel to estimate the conditional belief.

More challenges occur when recovering utility parameters in Stages 1 and 2, in which we have multiple learnings. When forming likelihood, we need not only to update learning for the current stage decision but also to form a belief about the probability of an outcome in later stages, which is updated following Bayesian learning, conditional on earlier stage outcome. Notably, the decision to take outside goods in the current stage not only prohibits learning for current stage attributes but also limits that for the next stage. Compared with a traditional Bayesian learning model with only one attribute-learning, our model requires a more sophisticated framework, as it needs to accommodate multiple learning. Compared with structural models with a steady policy function, our policy function is not stationary but, rather, updated with a Bayesian learning process, conditional on former decisions in sequence. This sequential and conditional decision process naturally leads to a forward simulation-based method for drawing decisions sequentially to form likelihood. Given the outcome variables for the last transaction, we update perceived attribute values for all three learning processes for each individual and update policy functions for the next period based on updated perceived attribute values. Given updated policy functions and the observed outcome for the current stage, we simulate outcome decisions for later stages for the next time period. We keep this loop through each time period. Similar to Erdem and Keane (1996), we numerically integrate perceived uncertainty by simulating many times for each individual and form a numerically integrated likelihood function.

Our model shows a good fit with the data. Because our dependent variable is a discrete-choice variable, we use McFadden's *pseudo* R^2 (McFadden 1974) as a reference of model fit. The results presented in Table 2.3 show that *pseudo* R^2s of the three stages are all above 0.2, which is interpreted as excellent fit, according to McFadden's discussion of behavioral travel modeling (Hensher and Stopher 1979). To visualize the goodness of fit, we plot the average probability of decisions over different individuals for a specific cumulative number of transactions in Figure 2.2, where the red dots indicate the nonparametric averages from our data, and the black ones represent predictive values from our parametric model. The *y*-axes in Figure 2.2's charts represent the probability that drivers accept orders, that orders are not canceled, and that passengers pay online for fulfilled orders. Nonparametric results in our charts show general ascending and concave patterns for all of those probabilities, which indicate that individuals undervalue TNC app-associated features at the beginning but gradually learn and use them more frequently with the accumulation of usage experience. By checking the fitness of the black dots, we find that our model recovers a learning pattern by a smooth concave and increasing function, which, in general, fits very well, except at the two tails, where the data are too sparse.

Table 2.3. Estimation Results

Parameter	**Estimate**	**St. Err.**
Decision to Accept an Order		
A_{a0}	-4.6644***	1.6006
A_{a1}	0	-fixed
σ^2_{a0}	10	-fixed
$\log(\sigma^2_{a1})$	1.7622	5.9296
$\log(\beta_{a1})$ (attribute)	-0.07910***	0.0275
β_{a2} (monetary scalar)	14.3350***	3.2362
c_{a2} (canceling cost)	12.0888***	1.1631
c_{a1} (online-pay cost)	-15.2681***	4.9458
c_{a0}	-8.3884**	0.8520
Pseudo R^2	0.5435	
Decision to Fulfill an Order		
A_{s0}	-1.3794***	0.0968
A_{s1}	0	-fixed
σ^2_{s0}	10	-fixed
$\log(\sigma^2_{s1})$	4.9423***	0.2209
$\log(\beta_{s1})$ (attribute)	-0.6848***	0.0705
β_{s2} (platform subsidy)	1.0418***	0.0455
β_{s3} (passenger tips)	-0.0990**	0.0431
β_{s4} (driver cashback)	7.1608***	0.1883
c_{s1} (online-pay cost)	-10.1917***	0.2217
c_{s0}	-1.3392***	0.0487
Pseudo R^2	0.3613	
Decision to Pay Online		
A_{p0}	-5.8569***	0.2959
A_{p1}	0	-fixed
σ^2_{p0}	10	-fixed
$\log(\sigma^2_{p1})$	3.2971***	0.0719
$\log(\beta_{p1})$ (attribute)	-0.4233***	0.0473
β_{p2} (passenger cashback)	0.2379***	0.0075
β_{p3} (passenger tips)	-1.2047***	0.0475
c_{p0}	- 1.3423***	0.0247
Pseudo R^2	0.2499	

Note: ** and *** denote significance at 5% and 1%, respectively.

Figure 2.2. Fitness per a Comparison of Nonparametric Aggregate Estimation vs. Prediction from Our Model

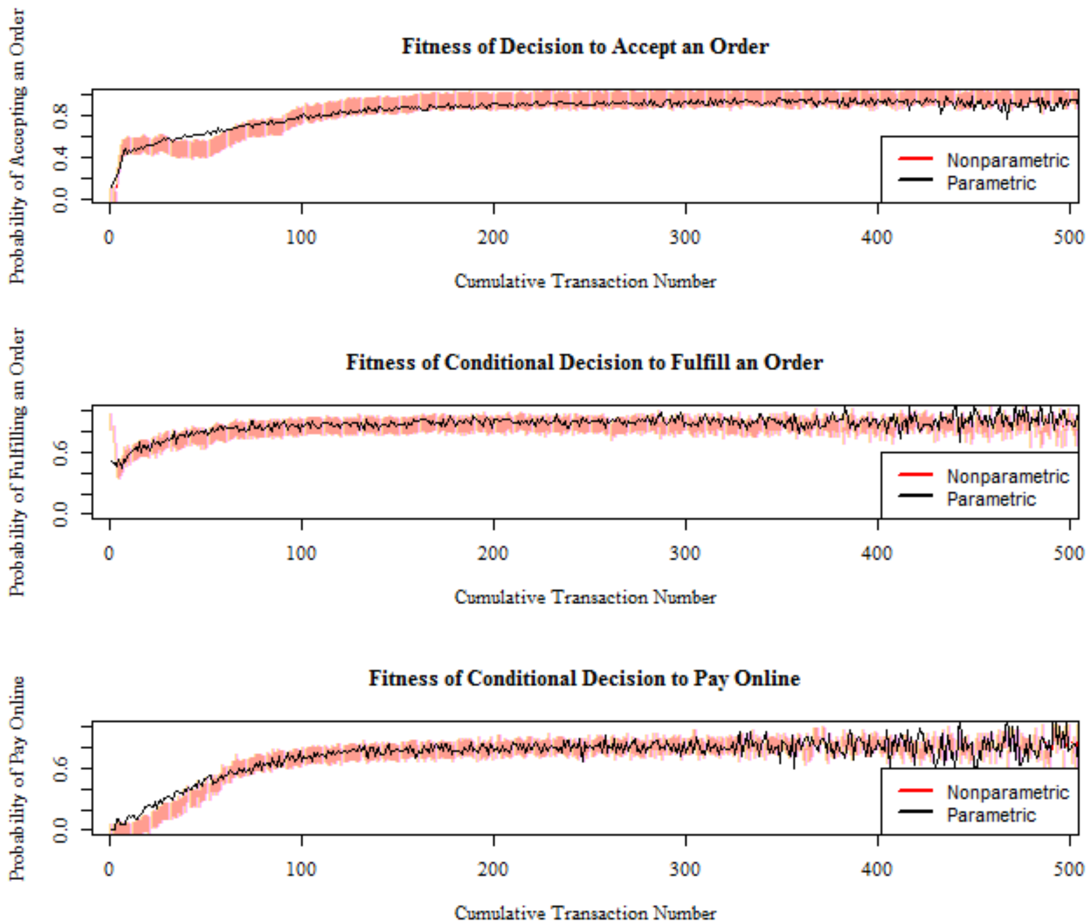

2.5.3 *Estimation Results*

In this section, we present the estimation results and insights associated with the estimates. We present the results in the same order as seen in our model, starting from the decisions of drivers and followed by the state transition of drivers' beliefs about passengers' decisions to use online-pay. All point estimates as well as standard errors are presented in Table 2.3. In addition, we show the latent learning process by explicitly displaying the updated perceived attributes and updated perceived variance parameters of the three decisions to help to understand the learning process. The attributes are simulated based on estimated parameters, which are shown in Figure 2.3.

Figure 2.3. Attribute Values and Uncertainties of Learning Processes

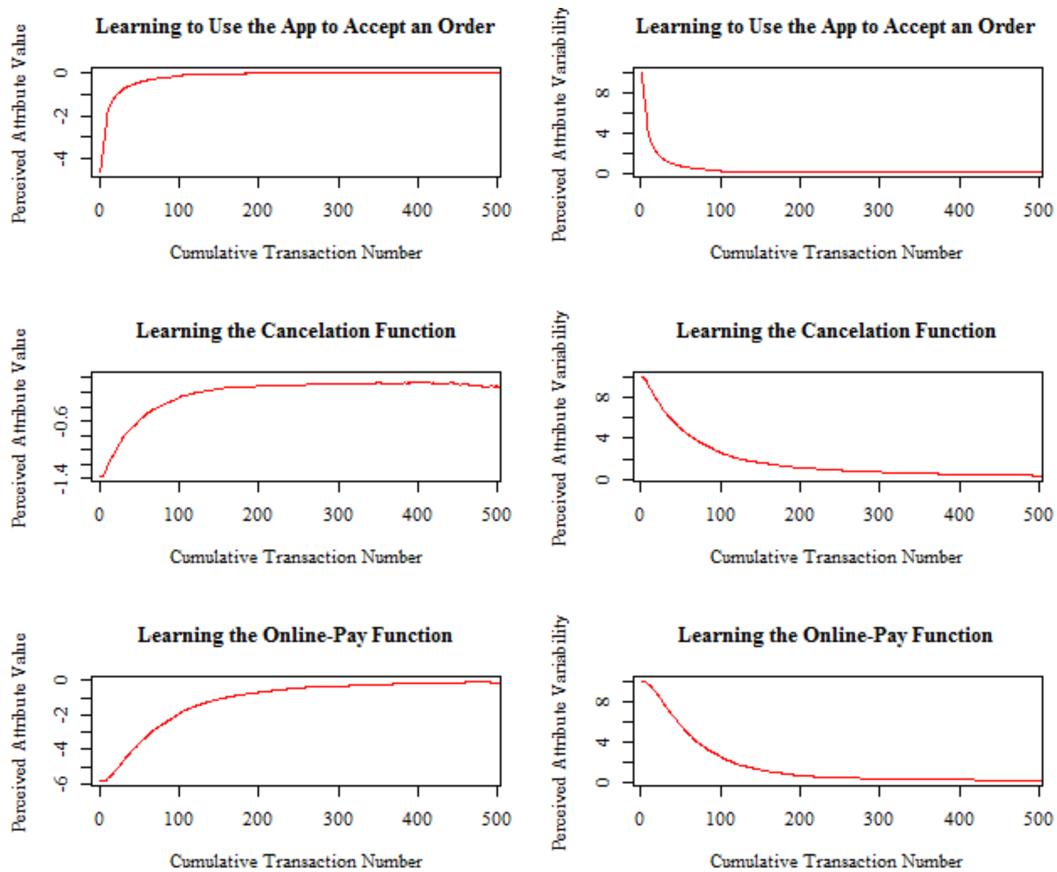

2.5.3.1 stage 1

As shown in the top panel of Table 2.3, we estimate the parameter of sensitivity to monetary rewards to be 14.3350, which implies the effectiveness of expected sales promotion in drivers' willingness to accept orders from the TNC channel. The expected cash back is due to drivers' beliefs about passengers' willingness to redeem the cashback and the probability of canceling the order in the next stage, which implies that drivers are more willing to accept an order when they believe that passengers are highly likely to redeem a sales promotion, the probability that they will cancel the order in the next stage is low, and a large amount in regard to the sales promotion is promised.

Our model allows us to recover latent utility parameters associated with the online-pay function, canceling function, and opportunity cost for taking outside goods. Specifically, the utility for canceling an order c_{a2} is estimated to be 12.0888, which is equivalent to 0.85 USD, once we adjust it with the monetary coefficient. This measure indicates that allowing strategic canceling leads drivers to be better off, as they can eliminate the orders that signal worse quality and higher risk. The cost of using online-pay c_{a1} is estimated to be -15.2681, or 1.07 USD. Such a cost might result from the transaction cost of using online-pay as well as the discounted utility due to delayed gratification. The opportunity cost for outside goods c_{a0} is -8.3884, implying that outside goods, in general, generate inferior attribute value when we have no online-pay option, disallow the cancelation feature, and have no monetary rewards for drivers. However, coupled with the cost of using online-pay, the utility of a canceled order, and associated probabilities, outside goods might show superiority. This is consistent with the intuition that usage complexity needs to be compensated. In regard to the effectiveness of monetary rewards as well as utility incurred from online-pay and canceling in the scenario when uncertainty is eliminated, a simple calculation from our results shows that compensation as low as 0.65 USD is enough to cover the cost incurred by using the app, which justifies the necessity of a subsidy for current TNC companies.

We also find strong evidence of learning. The prior parameter for app attribute A_{a0} is -4.6644, which is significantly lower than the normalized posterior parameter A_{a1}, which is normalized to 0. This difference indicates the undervaluation of the TNC app by drivers before they start to use it and suggests that learning can help to eliminate the perception bias gradually. We estimate $\log(\sigma_{a1}^2)$ to be 1.7622, or the variance of signal σ_{a1}^2 of 5.8252, smaller than the 10-normalized experience variability parameter. This indicates that the signal is much clearer than is the prior belief. Thus, the drivers' learning process in regard to the attribute of using the app to accept an order is comparably fast. The smaller the variability parameter, the greater the weight of the signal from usage experience and the faster the learning process is. Note that our estimation results in Stages 2 and 3 show that the learning processes for using online-pay and canceling orders are comparably slow, and our results demonstrate why the TNC platform in our example uses sales promotion associated with the online-pay function rather than with per-use.

2.5.3.2 Stage 2

The positive and significant coefficient on sales promotion for drivers from cash back indicates the effectiveness of the sales promotion in reducing drivers' tendency to cancel an accepted order. Specifically, as cash back from the platform is an expectation of drivers' belief about the probability of receiving cash back (which is the same as the probability of using online-pay), a positive coefficient suggests that the higher the perceived probability of receiving cash back, the less the drivers are willing to cancel an order. It is interesting to observe that, coupled with the claim in Stage 3 that the belief in the probability of receiving cash back is a function of a platform sales promotion strategy for passengers, the sales promotion strategy on the passenger side has an indirect and positive effect on drivers' willingness to fulfill an accepted order. In contrast, we estimate the coefficient of passengers' tips to be slightly negative (-0.0990). This supports our conjecture of two roles of passengers' tips. On the one hand, tips result in a higher utility level by providing a direct monetary incentive. On the other hand, tips signal the poor quality of a specific order such that passengers have to reveal tips to increase the probability that their orders could be taken, which consequently increases drivers' propensity to cancel. In our example, the second role of revealed tips dominates the first. In contrast, the positive sign of the coefficient on platform subsidy implies that subsidy either signals very limited information about quality or does not signal at all, leaving the monetary incentive to dominate the impact on utility function. This is consistent with our expectation, as the subsidy at an early stage is not dynamic but, rather, a constant for several short-term periods for all drivers and, thus, conveys very limited order-specific information.

The attribute prior parameter A_{s0} is estimated to be -1.3794 with posterior parameter A_{s1} normalized to 0, implying that drivers undervalue the perceived utility under uncertainty and, thus, overact by canceling existing orders. The experience variability parameter $\log(\sigma_{s1}^2)$ is estimated to be 4.9423. Compared with 10-normalized prior quality variance, it is much larger and, hence, displays a slow learning process. Similar to our argument about the direct versus indirect effects of sales promotion, this learning pattern confirms our previous conjecture that the platform has to sustain sales promotion for a certain length of time to ensure adequate learning.

We also identify two cost parameters associated with Stage 2. Given that the true attribute value is normalized to 0, the opportunity cost parameter c_{s0} is estimated to be -1.3392, showing an inferior utility for canceled orders. The cost associated with using online-pay is -10.1917, which is consistent with our finding in Stage 1.

2.5.3.3 Stage 3

From the drivers' perspective, passengers would be inclined to use online-pay when the amount of cash back from the platform, as a bonus for using online-pay, is high. This implies the effectiveness of the sales promotion on the passenger side (β_{p2} is estimated to be 0.2379). In addition, the higher the tips set by passengers, the lower is the willingness to pay online (-1.2047 coefficient estimate). This indicates that passengers who are generous with tips are less sensitive to a monetary sales promotion and, hence, less likely to redeem a cashback bonus by switching to an unfamiliar payment method.

With regard to learning, given that we fix the true attribute A_{p1} at 0, the estimated negative prior attribute A_{p0} indicates the undervaluation of using online-pay before drivers start to use and are informed about the quality of the app. In addition, to provide a comparison with the prior uncertainty of attribute σ_{p0}^2 fixed at 10, we estimate $\log(\sigma_{p1}^2)$ to be 3.2971; therefore σ_{p1}^2 is much higher than σ_{p0}^2, indicating a slow learning speed of the passenger population. The utility for outside goods is estimated to be -1.3423 with the attribute of online-pay's being normalized at 0. Thus, the online-pay function, in general, generates a positive utility for passengers compared with traditional cash transactions.

2.6 Policy Simulation

We conduct three sets of policy simulations to understand the impact of two-sided promotion on outcome variables. Specifically, we set our outcome variable as fulfillment rate, which is equivalent to the joint probability that an order is accepted and not canceled, and online-pay rate as the final outcome because they are the key measurements of TNC platform performance. In the first set of policy simulations, we focus on the time dimension, in which we attempt to differentiate the learning-related indirect effect from the sales promotion-induced direct effect on platform performance. The second set of simulations focuses on the effects of promotion on two sides. Our estimation results in the previous sections show qualitatively how two sides of the sales promotion affect drivers' willingness to use the app but do not tease out the effect of sales promotion from each side due to the indirect link among different stages. Therefore, our second set of policy simulations compares the effectiveness of sales promotion on two sides and generates insight into how to balance sales promotion between two sides with the objective of maximizing the overall platform performance. Given the findings from our model estimation and first two sets of policy simulations, in the third set of policy simulations, we modify cashback policies by following some heuristic rules, with a goal of finding a policy that increases online-pay rate while being more cost effective for the platform.

We use our estimation results to simulate a driver's decision to accept an order, whether a transaction is fulfilled, and whether a driver receives a cashback bonus. Specifically, based on our structural model, in each loop of simulation, we first use the updated belief of the probability that a passenger will use online-pay conditional on the transaction's being fulfilled and that the driver will fulfill the order conditional on accepting the order to simulate a driver's decision to accept an order. Then, we simulate whether the transaction is indeed fulfilled and whether the transaction is completed with online-pay. Finally, we update the perceived values of three attributes by using Bayesian learning and use the updated beliefs in the next time period. We simulate each policy 100 times and take the average as our result.

2.6.1 *Estimating Indirect Effect of Sales Promotion*

The indirect effect of sales promotion is implicitly shown by the existence of drivers' learning in our model. However, such an effect is difficult to measure directly by the estimated parameters. We conduct a set of simulations here to show explicitly how much impact such an effect has on drivers' decisions. We first simulate a baseline of drivers' decisions across time periods. Then, by identifying a benchmark point on a time horizon when the perceived value is stationary around the true value from that point on, or learning is no longer updated, we visualize the direct effect by simulating a case of 0.15 USD (equivalent to 1 RMB) cashback reduction in sales promotion policy from the benchmark point. Finally, we simulate the third case with a penalty before the benchmark point such that usage experience as well as learning is prohibited; then we set promotion to again be 0.15 USD less from the benchmark point. Because sales promotion is identical from the benchmark point such that the direct effect is the same, we visualize the indirect effect by presenting the difference between the second and third simulations, which is explained by the learning effect before the benchmark point only.

Figure 2.4. Direct and Indirect Effects of Sales Promotion

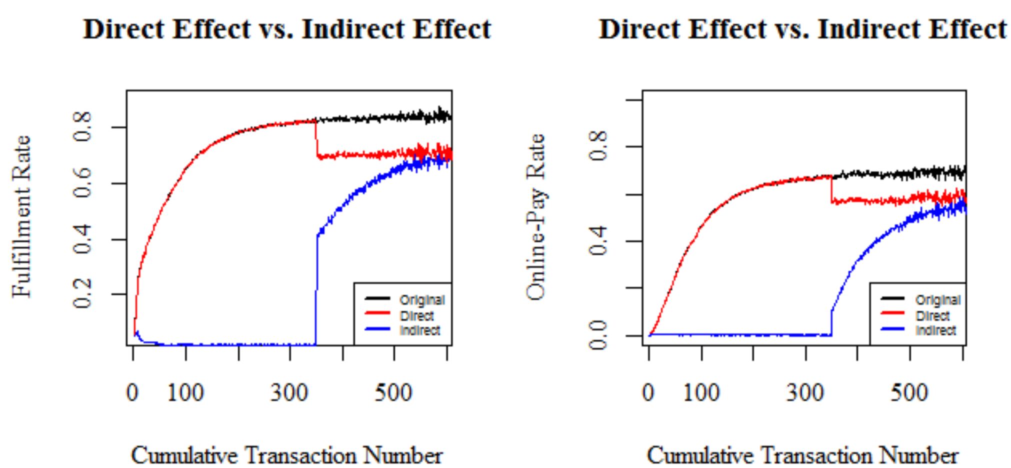

We aggregate the fulfillment rate and online-pay rate with respect to the cumulative transaction number and show our results in Figure 2.4. In the baseline case (black line), due to the fact that the sales promotion exists from a certain point until the end, both the direct and indirect effects exist at the end. We identify time point 350 (on the *x*-axis) as the benchmark point, as there is no more learning afterward, as shown in Figure 2.3. Therefore, we simulate the red line with a decrease of 0.15 USD from time point 350 to represent the case of the removal of the direct effect of 0.15 USD in the cashback promotion. The simulation result shows that fulfillment rate and online-pay rate are stationary from that point, in both black and red lines, validating that drivers' perceived attribute value is consistent with the true value and that the difference between the red line and baseline can be attributed solely to the direct effect of sales promotion. Finally, we simulate the blue line, which removes not only the direct effect of 0.15 USD but also the indirect effect from the sales promotion before the benchmark point. To achieve this, we impose a negative sales promotion of -5 USD before the benchmark point, which inhibits most of the order fulfillment and eliminates drivers' learning through usage experience. We apply the same sales promotion strategy (0.15 USD cashback reduction), seen in the blue line, after the benchmark and guarantee an identical direct effect of sales promotion for the blue and red cases. As a result, the difference between the blue and red cases after the benchmark point will be explained by past learning, which is the only difference between these two cases. Because the past learning is introduced by the past sales promotion policy, we can attribute the difference to the consequence of the indirect effect of sales promotion in the earlier time period. As we can observe, the blue line still follows an increasing pattern, implying that drivers are still learning about the true value of the attributes of the app. This indirect effect will diminish with the accumulation of usage experience.

2.6.2 *Estimating the Effect of Sales Promotion on Passenger Side*

Our model estimation results show that sales promotion on the passenger side has an impact on outcome variables. However, the effect itself is not explicit, given the non-linear function form, casting doubt on how strong such an effect compares with the effect of sales promotion on the driver side. We simulate the outcome variables by adjusting passenger side sales promotion policy to visualize this effect explicitly. In addition, we simulate the outcome variables with an identical adjustment to the driver side sales promotion to comparatively illustrate the strength of passenger side sales promotion effect.

We first simulate the outcome variables of fulfillment rate and online-pay rate with the original sales promotion policy, shown as the black line. Given an adjustment of reducing 0.15 USD (1 RMB) on passenger side cash back, we simulate the outcome variables represented by the blue line. Further, we apply the same adjustment to the driver side, shown as the red line. The graph on the left in

Figure 2.5 shows that, though the present, the marginal effect of the passenger-side sales promotion is very limited in terms of the fulfillment rate compared with that of the driver-side sales promotion. With respect to online-pay rate (the graph on the right), as sales promotion on the passenger side has a direct effect on the use of online-pay, its overall marginal effect is more significant but still small compared with that of the driver side.

Figure 2.5. Decrease in Passenger Side Sales Promotion by 0.15 USD

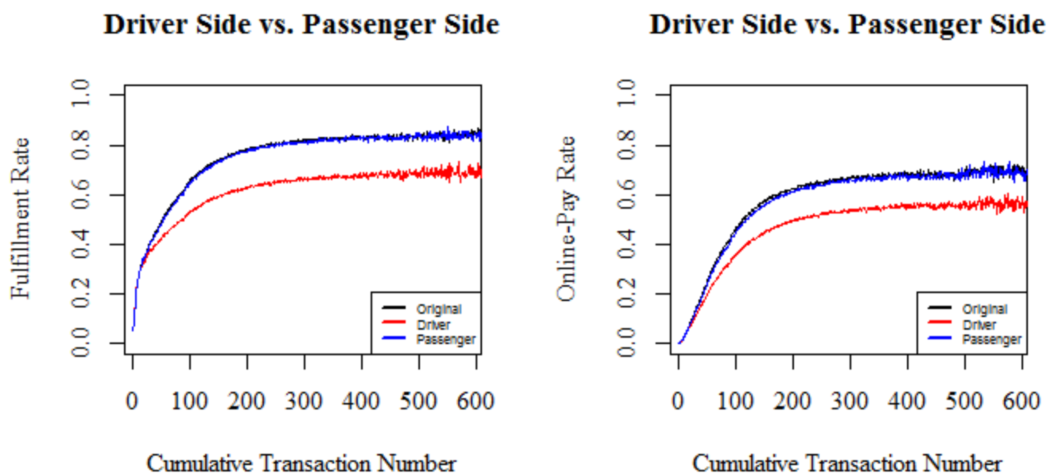

Figure 2.6. Decrease in Passenger Side Sales Promotion by 1 USD

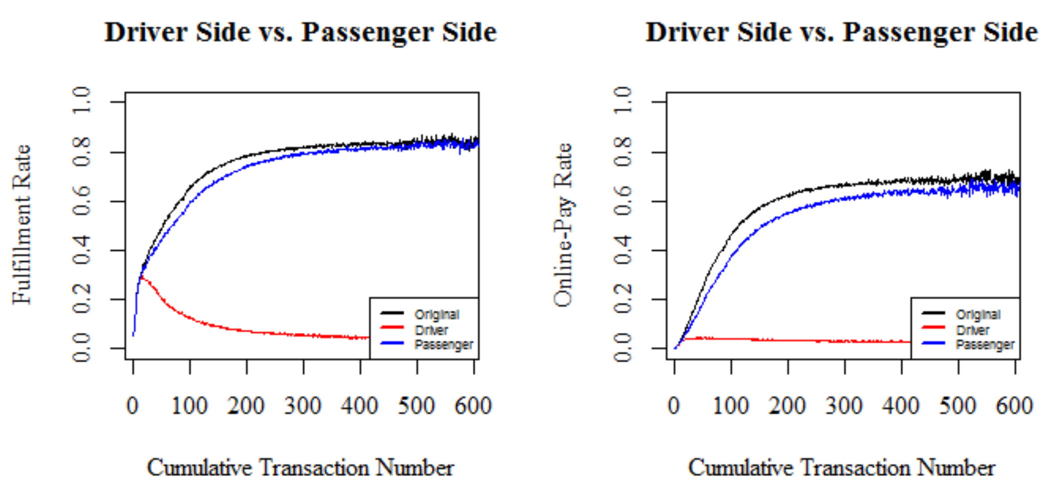

To ensure the robustness of our findings, we specify a more powerful adjustment toward sales promotion to test whether similar results will be maintained. In our experiment, represented in

Figure 2.6, we reduce the sales promotion by 1 USD on either passenger side or driver side. Compared with our findings in

Figure 2.5, the marginal effects of sales promotion adjustment are more significant for both fulfillment rate and online-pay rate, and the conclusion that driver-side effect is more prominent still holds in the new experiment. Note that our research is conditional on fixing order-generating rate due to data limitations; this observation leads to a managerial suggestion that the TNC should put more weight on the driver side to optimize the overall performance of the app, conditional on the fixed arrival of orders.

2.6.3 *Optimizing Sales Promotion*

In this section, we apply model estimation results and earlier policy simulations to optimize the sales promotion strategy for the TNC platform. Here we do not attempt to manipulate the policies to achieve the optimality, but, rather, we show directive heuristic policymaking methods, which, by following, we can improve the platform performance. Our objective consists of improving two dimensions: online-pay rate and sales promotion cost. We focus on online-pay rate because it is our final outcome, following the three-stage sequential decisions, which reflects the overall impact of sales promotion. In addition, it is also the only stage associated with monetary transactions and revenue inflow. A TNC that increases overall online-pay rate and reduces sales promotion cost will be better off with respect to profit maximization. The heuristic policy that we use is based on our findings of indirect effect generated from learning and the comparison between the effects of sales promotion on two sides. Because sales promotion of the passenger side shows a weaker effect compared with that of the driver side, our first adjustment is to rebalance sales promotion between the two sides with more weight on the driver side. Further, we rebalance sales promotion with regard to time. Specifically, we use intensive sales promotion at the beginning of the product introduction period to achieve a more indirect effect. This is consistent with our claim that sales promotion has a marginal indirect effect when learning happens.

Again, we start by simulating a baseline by following the original policy, shown as the black line in Figure 2.7. In this policy, sales promotion exists for both sides, with an average of 1.47 USD for the driver side and 1.58 USD for the passenger side and an overall sales promotion cost of 285,086 USD for 952 drivers. The first adjustment we make is to remove all sales promotion on the passenger side and increase the sales promotion on the driver side by 0.1 USD, shown as the red line. The overall sales promotion cost decreases significantly to 116,610 USD, along with a significant loss of overall online-pay usage, shown as the space between the black and red lines. The online-pay rates of the black and red lines converge to a similar level with the accumulation of usage experience, implying similar direct effects of these two sales promotion policies. The difference between the two lines is explained mainly by an inadequate indirect effect at the earlier stage.

Figure 2.7. Online-Pay Rate with Improved Sales Promotion Strategy

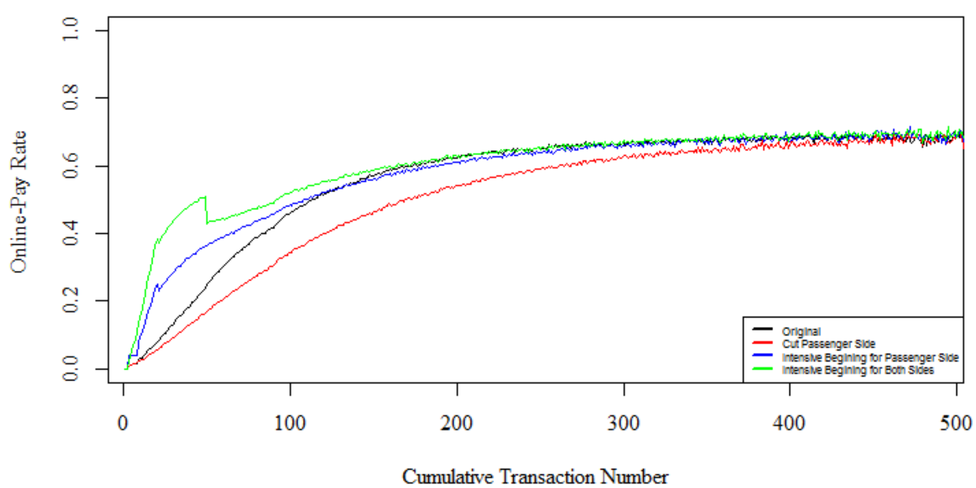

To alleviate the loss due to inadequate earlier-stage sales promotion, we increase sales promotion on the passenger side to 4 USD during the first month after the app is released, reducing the overall cost to 186,010 USD. From the blue line, we find that intensive sales promotion works effectively to increase the overall online-pay rate by both direct and indirect effects, shown as the space between the blue and red lines. The effects occur not only on the left-hand side of the figure, when the increased sales promotion is imposed, but also in the regions on the right-hand side, where an indirect effect through learning remains. The online-pay rate, however, still falls below the original black line after 150 transactions due to slower learning that results from insufficient incentivizing.

Our final policy is based on the blue line, with a policy change of sales promotion of 3 USD for the first 50 orders on the driver side. The intensity of this short-period promotion doubles the average in the original case. The reason that we offer only 50 orders with a doubled sales promotion is that learning to accept an order and learning to cancel an order are comparably fast. As seen in Figure 2.3, we find that learning to accept orders finishes with 30 orders and that learning to cancel orders finishes with 100 orders. A more intensive sales promotion will lead to an even faster learning process, and, therefore, we choose 50 orders as the duration. The green line in Figure 2.7 represents the overall performance of this policy, with the space between green and blue lines as representing the marginal effect of an early sales promotion on the driver side. The figure also shows the indirect effect through the gap between blue and green lines after a cumulative transaction number of 50. Our final policy envelops the original one (black line), indicating an overall improvement with respect to online-pay rate. In addition, the overall cost is 238,338 USD, much less than 285,086 USD in the original policy and, consequently, results in higher platform profit. This finding explains why many TNCs have an intensive sales promotion at the very beginning.

It is worth noting that, even when learning is improved by imposing an intensive sales promotion, this does not necessarily imply that such a promotion should last until the learning finishes. In fact, none of our simulations leads to completed learning by the end of an intensive sales promotion. Although learning can lead to higher perceived attribute value, its marginal effect diminishes with the accumulation of usage experience, and there will be a point where the marginal effect of learning is inadequate to cover the cost incurred by sales promotion.

2.7 CONCLUSION AND IMPLICATIONS

TNC app platforms are attempting to reduce their reliance on sales promotion to gain new users, especially drivers. In our research, we find a means to reduce this reliance by developing an understanding of two mechanisms: first, how new users form their initial preference of TNC apps; and, second, how sales promotion affects the formation of initial learning.

Our analysis shows that individuals are conservative about adopting a new channel. TNC app users underestimate not only the attribute values of using an app but also the perceived attribute value of passengers' preference to use an online-pay function. As a result, drivers form a comparably lower willingness to use the app and higher willingness to cancel a potentially "bad" order due to lower prior attributes and a lower probability of receiving a cashback bonus. This is consistent with our preliminary results that drivers start with a lower frequency of using the app functions, whereas the canceling rate is higher.

Usage experience plays a significant role in alleviating the bias from uncertainty and risk aversion. With the accumulation of usage experience, drivers obtain adequate exposure to the true attribute values of the app and consistently perceive the belief of passengers' decision to use online-pay with diminishing uncertainty. The more usage experience that drivers have, the faster their perceptions converge on the true attribute values. Given the estimation results, we know that drivers are more tolerant of accepting orders and fulfilling orders and are more optimistic about being rewarded with a cashback bonus once their learning process is completed. This finding supports the common industry practice of enhancing usage experience during product introduction to help consumer learning and justifies the effort of enhancing usage experience through different marketing tools, including a sales promotion, which is the example in our research context.

Our counterfactual analyses confirm sales promotion as an effective tool for user learning. As a component in user utility function, sales promotion affects drivers' decisions to accept an order at any time. Our results show that a sales promotion in the introductory period has a more significant effect than a promotion that is introduced later. Other than the direct effect of sales promotion on contemporaneous drivers' utility, an early sales promotion affects decisions in later periods through early enhanced usage experience due to user learning. This finding justifies the importance of early sales promotion for a recently introduced product and explains why most TNC app platforms put forth enormous effort toward sales promotion in early stages. In addition, we observe the cross-effect of two-sided sales promotion. Sales promotion for passengers can influence the decisions of drivers when drivers are rationally forward-looking and the decisions of passengers affect drivers' earnings. We further propose a heuristic method that takes advantage of both direct and indirect effects to help app platforms increase their income while controlling their cost for early sales promotion. In addition, we examine the performance improvement due to rebalancing sales promotion for two sides by lowering the weight for the less effective passenger side.

Our paper has several limitations. First, the data available for researchers are limited. The introductory period of the product might indicate incomprehensiveness of data maintenance. For each transaction, our data do not include such information as pickup locations, destinations, duration, or traffic routes, even though drivers might be able to partially obtain or infer such information. Such a limitation restricts our investigation under the assumption of homogeneity of orders in terms of drivers' learning process. Second, we take sales promotion as exogenous. Research that considers strategic sales promotion might could allow investigation of competition among different platforms and extend our learning framework to that of forward-looking. Such research requires a dataset with a long enough period of sufficient variance of sales promotion in equilibrium. Our data include only a period with several constant sales promotion amounts, each of which lasts for a long period. Individuals are well informed that, in the coming period, the sales promotion is the same for all orders. In addition, the short introductory period may not support the equilibrium assumption. Third, we do not account for network effects in our model. Network effects can be one component of learning signals, such that they may have an impact on the speed of learning. Due to the sparsity of observation and limited information, however, it is very difficult for us to determine the "learning from network" mechanism and recover the network structure. All of these limitations can be addressed with additional data.

Despite the limitations, our paper makes the following contributions. First, it is the first study that econometrically models drivers' decision process in the use of a TNC app. Our model captures how drivers' decisions are influenced by TNC monetary rewards as well as their perceived passengers' decision. In addition to the direct effect from two-sided sales promotion, we also depict drivers' learning from their usage experience that indirectly contributes to their overall use of the app. Our results describe drivers' learning of multiple attributes of the TNC app. Second, we run policy simulations to examine differential effects from different sets of marketing promotion designs. This generates managerial insights for the runners of newly-introduced products about designing a sales promotion in a more effective way. Our counterfactual analyses suggest that an early, intensive sales promotion policy not only enhances users' willingness to use but is also cost effective.

Chapter 3. INFORMATION ASYMMETRY IN ONLINE TWO-SIDED MARKET - ONLINE SERVICE MARKETPLACE

3.1 INTRODUCTION

The trend toward digitalization has expanded the product line of the online marketplace from traditional physical products to services. A service marketplace connects businesses with consumers who need to do such things as purchase or recharge a prepaid phone card or find a babysitter, expanding greatly the breadth of and demand for the online marketplace. This expansion makes the marketplace more transparent, easier to use, and more popular with online consumers. In the case of prepaid phone card purchasing, for example, without an online tool, consumers have to find and contact retailers individually, compare prices, and then make a decision. In the online marketplace, these steps are completed in several minutes through the use of key words and filters. Services vary from simple home improvement tasks to more sophisticated ones, such as coaching sessions, professional editing, and even system development. Well-funded marketplaces include home improvement platforms, such as Thumbtack, Amazon services, Handy, and Zaarly; freelance platforms, such as Elance, Freelancer, and ODesk; and more specialized platforms, such as Skillbridge for business consulting, Vouched for financial services, and UpCounsel for legal services. Industry observers estimate a U.S. market size of at least $250 billion for the home repair and improvement market alone (Seetharaman 2014). Other reports indicate a profit margin of about 20% for the service marketplace by Amazon, suggesting the potential for huge profitability (Perez and Etherington 2015).

Compared with physical items sold in the online marketplace, service is more vulnerable to information asymmetry. Because service is intangible and insubstantial, consumers can find it more difficult to accurately gauge the quality of service when comparing measurable characteristics among alternatives as compared to the process as applied to physical products. In addition, each service is generated one time and is unique, which results in variability and inconsistency of services, even by the same provider. Consumers might need additional information (e.g., a larger sample of user experience) to develop a more accurate expectation of the quality of a service and a risk-hedging tool to alleviate their uncertainty. This makes information provided by the platform more crucial to the purchase decision-making process of consumers and, thus, the demand for sellers.

A platform provides three forms of information, categorized by the identities of its generators. The first form is information provided solely by sellers, such as price, product variety, and length of membership in the platform. This classic category of information provides fundamental information about price and availability to help sellers compete for a consumer. This category exists widely not only in the service online marketplace but also in traditional marketplaces and has been well addressed in past research.

The second category is consumer-generated reputation. This information is created directly or indirectly and is derived from consumer evaluations once transactions are fulfilled and includes, for example, rating valence and rating volume. Typically, the platform manager verifies the purchases and user reviews, leading to the credibility and robustness of consumer-generated reputation, as no seller can easily manipulate his or her own reviews or ratings or take down a competitor by malicious tampering. The reputation of a seller mitigates consumers' uncertainty associated with the service by providing a predictive distribution of service quality inferred from past transactions. Reputation informs consumers of expected means of satisfaction, product variability denoted by variance, and popularity, based on cumulative evaluations volume. The effectiveness of reputation in mitigating information asymmetry in the online service marketplace has been rigorously examined and validated (e.g., Yoganarasimhan 2013).

The consumer is not the only contributor to a seller's reputation-related information. As a complement, coordination between platforms and sellers contributes to the third type of information—platform endorsement. Shown as unified logos or tags, platform endorsement is widely applied in most online marketplaces; examples include "trusted seller" certifications or "on sale" tags. Because quality of service is difficult to measure due to the lack of unified characteristics of marketplace spaces, platform endorsement provides sellers with a platform-wide standardized characteristic space, allowing consumers to evaluate and compare different sellers within a measurable system. The availability of platform endorsement benefits sellers and the platform because, different from consumer-generated reputation, platform endorsement can be implemented, adjusted, and monitored easily by sellers and by the platform. Therefore, sellers can make use of endorsement as a strategic tool, in addition to price, to increase competency, whereas platforms can apply an endorsement as a tool for seller regulation and, thus, improve platform-wide quality and the ability to compete with other marketplaces.

Although platform endorsement is widely used, there is no consensus on its effectiveness due to a number of issues. First, unlike consumer-generated reputation, which captures an overall measurement of service quality, each platform endorsement focuses on only one specific characteristic, such as accepting coupons or having platform refund insurance. A certain characteristic might be favored by a type of consumer who values that characteristic more than others. As such, consumers on the platform might have heterogeneous tastes for platform endorsement. In addition, "There's no such thing as a free lunch," which, in this case means that platform endorsement is not free for sellers who use it. A rational consumer would be able to infer that the cost of platform endorsement can be found in the markup in the price paid by consumers. For example, when the endorsement, logo, or tag concerns gift giving, only consumers who favor that gift would be willing to buy the service or product. For other consumers, the endorsement, logo, or tag might be redundant and signal an additional price that they would have to pay for the gift. They thus switch to other sellers without such an endorsement. This example demonstrates the possibility that consumers might perceive a platform endorsement as negative, even though they are indifferent to the endorsement itself, implying the importance of taking into account the heterogeneity of consumers' tastes. Further, unlike consumer-generated reputation, which can be treated as exogenous, platform endorsements are potentially endogenous, as sellers have the freedom to easily make strategic adjustments based on the unobserved shock of demand. The correlation to be reckoned between unobserved characteristics and platform endorsement constitutes a classic endogeneity problem similar to that for prices. An examination of the effectiveness of platform endorsement will be biased unless the endogeneity issue is addressed.

The first objective of this paper is to examine the effectiveness of platform endorsement and consumer-generated reputation on demand in the online service marketplace. Specifically, we apply a demand estimation method that uses an aggregate data (BLP; Berry et al. 1995) model to quantify how platform endorsement, consumer-generated reputation, and other seller-specific information affect consumers' willingness to purchase and, thus, demand for sellers. The BLP model incorporates random coefficients to allow heterogeneous taste coefficients across consumers and instrument variables to provide unbiased estimates with endogeneity taken into consideration. We apply the model to the data for the specialized service of selling prepaid phone cards and recharging them remotely. The model allows us to derive numerical marginal effects and elasticities for a consumer-generated reputation system and the counterfactual marginal effects of platform endorsement to understand the effectiveness of available information from the seller's perspective. The model also enables closed-form consumer surplus to address the consumer-side impact of the platform endorsement policy and consumer-generated reputation.

During our observation, the platform announces a "conform or be cast out" policy to encourage sellers' registration in a platform refund insurance program as one specific type of platform endorsement. A platform refund insurance is used to alleviate consumers' risk aversion to online service by providing a platform-wide insurance that guarantees a full refund of the paid amount in the form of a credit on the platform if consumers dispute a transaction and request a refund. The program benefits consumers by improving the after-sale warranty; however, sellers incur a cost in the form of an insurance premium, resulting in an overall low percentage of insured sellers. To improve the platform-wide competency and service quality, the platform put into effect a "conform or be cast out" policy to reach a goal of 100% platform refund-insured sellers in the marketplace. In particular, the platform set Week 10 as the start point of the policy and cast out unregistered sellers in the following four weeks. Unregistered sellers who want to continue their business on the platform have to conform by joining the program before being expelled.

The "conform or be cast out" policy is widely applied by online platform management. It is straightforward to implement by the platform, easy to understand by sellers, and clear in regard to the policy orientation and objective, endowing it with popularity among platform managers. A theory-based understanding suggests multisided effects in the "conform or be cast out" policy for different stakeholders. For consumers and the platform, the policy comes with a positive aspect, as an improved rate of insured sellers, and negative one, as "casting out" results in a decreased variety of sellers. The "casting out" rule rewards sellers who decide to stay by squeezing demand from the market share of expelled sellers, whereas the higher rate of insured sellers mitigates the competitive advantage of sellers who adopted the program before the policy. The coexistence of positive and negative effects of the policy for relevant stakeholders leads to further uncertainty about the value of the policy, suggesting the importance of quantitatively measuring distinct aspects of policy effects and potential indirect effects. To our knowledge, however, there is no prior literature that empirically quantifies the effects of the "conform or be cast out" policy in the online marketplace.

To address this research gap, our second objective is to analyze the mechanism and effects of the "conform or be cast out" policy. Note that the policy consists of two components: "conform" and "cast out," which exhibit distinct, direct effects. Descriptive demand analysis identifies only the compound direct effect for sellers and the platform, suggesting the need to use a model-based approach to decompose direct effect into distinct components and to understand how they affect consumer welfare, correspondingly. Further, the exogenous policy changes the market structure, possibly leading to a secondary strategic reaction of sellers and establishing a new equilibrium. The impact of the secondary strategic reaction after the policy shock leads to an indirect effect of the policy. By using counterfactual simulation based on the estimated BLP model, we examine the direct effects of each side of the policy and compare equilibrium before and after the policy to infer potential indirect effects explicitly.

As we expected, our results show diversified preferences to several platform endorsements and more consistent sensitivities to consumer-generated reputation across consumers. With regard to consumer-generated reputation, faster service rates, fewer disputes, fewer required refunds and fulfilled refunds, higher average ratings, larger rating volume, and more positive evaluations and fewer negative ones increase demand for sellers. Except for the number of fulfilled refunds, all of the dimensions of consumer-generated reputation have relatively small variability parameters, which generate limited impact on demand. Taste coefficients for platform endorsement, however, are distributed dispersedly across consumers, except those for short-term sale discounts, platform refund insurance, VIP stores, and detailed pictures. Variability of tastes has a significant impact on the demand for particular sellers, resulting in a seller-specific marginal effect and a substitution effect of platform endorsement. When sellers adjust their platform endorsements, for example, through coupons, short-term VIP seller tags, threshold discounts, sample gifts, acceptance of credit cards, and guaranteed return and exchange, a strategy that increases demand for Seller A might have the opposite effect on the demand for Seller B. In addition, we find that platform endorsements such as seasonal short-term sale discounts generate a mainly negative effect on demand, consistent with our conjecture that an inferred markup affects purchasing behavior. Consistent with our intuition, our analysis shows that consumers prefer sellers with lower prices, more product variety, and a longer membership history in the platform. To understand the demand system from the perspective of sellers, we further calculate self- and cross-market share elasticities of reported delivery time and of price, showing the effectiveness of using price and a consumer-generated reputation system as tools for demand competition. To measure the impact of policy change on consumers and the platform, we derive consumer surplus from each market and find evidence of welfare loss due to decreased variety and welfare gain due to an increased rate of insured sellers.

With regard to the “conform or be cast out” policy, we decompose the effect of conforming to the implementation of platform refund insurance from that of being cast out if resisting. A simulated marginal analysis with adjustment by one seller and a counterfactual analysis that mimics the market structure change that we observed in data show that switching to platform refund insurance increases the demand for those sellers and decreases the demand for other sellers who are already in the platform refund insurance program. Coupled with the substitution for expelled sellers, the benefit for sellers who newly switch is prominent, and the loss of demand for early adopters due to increased competition will be offset fully or partially. From a consumer's perspective, we find that the welfare derived from an increased rate of insured sellers overwhelms the loss from less product variety, indicating an overall improvement of performance and competency of the platform. Finally, by comparing the consumer surplus of a market after the policy change with that of a counterfactual market that has the same level of platform refund insurance registry and the same number of sellers, with the maintenance of other characteristics in a market before the policy change, we find higher consumer surplus for a market with after-policy characteristics, indicating that sellers' reaction to the policy leads to further improvement in other characteristics. Despite no analytical solution for the formation of the new equilibrium, our finding suggests a positive indirect effect of the "conform or be cast out" policy on consumers and the platform.

The rest of the paper is organized as follows. In Section 3.2, we provide a brief discussion of the related literature and our contribution. We then present the research context and available information from our dataset in Section 3.3. In Section 3.4, we present demand estimation by following an aggregate random coefficient logit model (BLP model; Berry et al. 1995). In Section 3.5, we discuss our estimation strategy and explain how the endogeneity issue is rigorously controlled. We then report estimation results in Section 3.6, with elasticities derived, to understand competition among sellers, and consumer surplus recovered, to measure the impacts of policy on consumers. In Section 3.7, we simulate data to measure the direct and indirect effects of the "conform or be cast out" policy for platform refund insurance. Finally, we summarize our finding and conclude our research in Section 3.8.

3.2 LITERATURE REVIEW

Our paper builds on and extends the literature that examines how a reputation system, platform endorsement, and more generalized website design affect transactional outcomes in the context of the online marketplace. Researchers have addressed the effects of specific components of the reputation system or general reputation on different forms of outcomes, such as the duration of time to sell products (Ghose 2009) and an auction outcome (Bockstedt and Goh 2011) in the online goods marketplace. The components of reputation include positive numerical ratings (Ba and Pavlou 2002) and positive textual feedback comments (Pavlou and Dimoka 2006). Further, other information signals, such as diagnostic product descriptions and third-party product assurances (Dimoka et al. 2012) and the quality of e-images (Gregg and Walczak 2008), are shown to increase the price premium. Schlosser et al. (2006) find that website investment signals trustworthiness and increases consumers' intention to purchase. Wells et al. (2011) discuss how quality of the website affects impulsiveness of consumption.

To our knowledge, one of the closest studies, with respect to our research objective, was conducted by Li et al. (2009), who demonstrate that revealing quality and credibility indicators, such as ratings, a money-back guarantee, and third-party payment method, encourages bidders to participate in auctions. The closest research, with respect to context, was done by Yoganarasimhan (2013), who uses a structural model to measure the effectiveness of a reputation system in a typical service marketplace that is understood as a freelance marketplace. Both papers focus on the setting of online auctions. We differentiate our work from prior studies as follows: (1) we extend the information signal to platform endorsement, the effectiveness of which is challenging to measure due to its endogeneity; (2) we quantify the effects of the perspectives of all stakeholders, including the seller, consumer, and platform; (3) we focus on the outcome as a demand system to generate economics insights, such as competition; and (4) we address the heterogeneity of consumers' tastes in regard to the information displayed on websites.

Our paper also contributes to literature that quantifies consumer welfare in the online marketplace. To our knowledge, there are only a few papers that focus on measuring consumer welfare. Using a field experiment, Bapna et al. (2008) quantify a median surplus of at least $4 per eBay auction extracted by a consumer, suggesting, in 2003, a total surplus of at least $7.05 billion. Brynjolfsson et al. (2003) measure the consumer welfare of product variety in the context of the Amazon book store. They estimate the enhanced consumer welfare to be between $731 million and $1.03 billion in the year 2000, which is between 7 and 10 times that from increased competition and lower prices in the market. Similar to our method, Ghose and Han (2014) use the BLP model to estimate enhanced consumer surplus, gained from the availability of mobile apps, to be approximately $33.6 billion annually in the United States. We extend this stream of literature by measuring consumer welfare in the online service marketplace. In addition, using a simulation method, our work is the first to measure the direct and indirect utility changes of the "conform or be cast out" policy by calculating compensating variation.

Our research applies a demand estimation method, the well-known BLP model, that uses aggregate data (Berry et al. 1995). The model shows its superiority in demand estimation by allowing heterogeneity in consumer taste and endogeneity of product characteristics with only aggregate-level information about market structure. It also allows the estimation of consumer welfare, given its structure of individual-level decision making that maximizes latent utility. Researchers continue to develop the BLP model to enhance its applicability and performance, for example, by using consumer demographic information to improve precision (Petrin 2002) and applying an MPEC algorithm to accelerate the computation speed (Dube et al 2012). The model allows wide application across the literatures in empirical industrial organizations, marketing and information systems in a variety of industries, including the automobile industry (Berry et al. 1993, 1999, 2004; Petrin 2002), ready-to-eat cereal (Nevo 2001), movies (David 2001), online hotel booking (Ghose et al. 2012), and mobile app usage (Ghose and Han 2014). Our paper extends the applicability of the BLP model to the online service marketplace, which is characterized by high levels of information asymmetry and competitiveness. We also extend the applications of the BLP model to the context of large numbers of players and large spaces of endogenous actions, which are common challenges in research on the online marketplace.

3.3 Data and Descriptive Findings

Our data are from one of the world's largest online marketplaces based in China. The marketplace is famous for traditional physical goods but has been expanding business into service recently, with the addition of sites for travel agents, electronic device repairs, and prepaid phone recharging. We focus on online prepaid phone recharging because this submarket involves clear physical goods (prepaid cards) and a standardized service with a clear procedure; this avoids the inflation of the error term that captures unobserved characteristics of the service and leaves sellers' information available online as the main predictor of demand among sellers. This setup greatly simplifies the model's challenges and provides us with a good opportunity to investigate the formation of demand with available information online.

The market has sufficient size and a sufficient variety of sellers' characteristics for us to empirically identify the effectiveness of consumer-generated reputation and platform endorsement. Sohu IT (2014) found that prepaid phone cards occupied more than 50% of the market share of the telecommunication market in China, constituting enormous demand for recharging consumers' prepaid phones. As seen from the supply side, the market is fragmented unevenly. On any given day, there are approximately 20,000 active sellers available in the platform, with each megaseller's occupying as large as 5% of the total market, which is approximately 70,000 times that of a 25% quantile small seller. Note that providing effective information would alleviate information asymmetry in the online market, which is due to the intangibility and variability of services. The existence of heterogeneous seller size usually indicates the effectiveness of information, which helps consumers to avoid purchasing a "lemon" and leaves low-quality sellers behind. We expect the huge variance in demand to be explained by the reputation system and platform endorsement, as they are the major sources of information on which consumers base their purchase decisions.

We collect data for 48 weeks, with each week specified as one market. In each week, aggregate-level information on the seller is collected. Demand is measured as total number of sales for each week. Consumer-observable pivotal information for making decisions is collected regarding consumer-generated reputation, platform endorsement, and other seller-generated information, as summarized in Table 3.4. With regard to seller-generated reputation, we collect typical aggregated review information, such as rating valence and other consumer-reported information, that signals quality of service. Given that most consumer-reported information is managed as count data, such as number of disputes and number of refunds, we take a log transformation of the data, assuming a decreasing marginal effect when we apply the data in the model. All possible platform endorsements during the observation window are collected for each seller as dummy variables, with 1 as indicating the seller's state in terms of the endorsement. Whenever a seller registers a specific endorsement, the platform will attach a standardized logo to the seller on the search results page to highlight the information that sellers would like to present to consumers. Most platform endorsements in our study signal sellers' characteristics, such as their being VIP sellers, having sales promotions and extended warranties, and their products' being easy to use. By excluding sellers who have no transactions, we have a sample with 878,665 observations and 18,306 sellers per week, on average.

Table 3.4. Definition and Summary Statistics of Variables

Variable	Definition	Summary Statistics			
		Mean	Std. dev.	Min	Max
Consumer-generated Reputation					
Delivery time	Average of Reported delivery time in past transactions (day)	0.007	0.622	0.000	27.83
#Dispute	Cumulative count of disputes in past transactions	0.197	4.009	0	499
#Refund RQST	Cumulative count of required refund in past transactions	9.170	2.512	0	41141
AVG rating	Cumulative Average rating in past transactions	4.445	1.529	0	5
Rating volume	Cumulative Rating volume in past transactions	62.55	1169	0	455302
#Refund	Cumulative count of fulfilled refund in past transactions	0.053	1.464	0	300
#Pos EVAL	Cumulative count of transaction evaluated as positive	953.9	1.4E4	0	1505258
#Neg EVAL	Cumulative count of transaction evaluated as negative	5.383	56.48	0	6475
Platform Endorsement					
Coupon	Seller accept using coupon	0.009	0.092	0	1
Short-term VIP	Seller is periodic VIP seller	0.017	0.131	0	1
TD	Discount when threshold are met	0.003	0.055	0	1
Sample gifts	Giving gift	0.002	0.039	0	1
Short term sale	Flash sale	0.012	0.112	0	1
VIP store	Seller is VIP store	0.013	1.122	0	1
Credit card	Seller accept credit card	0.082	0.274	0	1
Detailed picture	Detailed picture of product available	2E-4	0.017	0	1
PRI	Platform refund insurance	0.715	0.451	0	1
SGR	Seller guarantee to refund with return	0.046	0.210	0	1
Seller-generated Information					
Product variety	Number of different products for a seller	394.9	254.6	1	15690
Duration	Length of seller membership on the platform (day)	342.6	434.3	7	3226
Price	Average reserve price during the week (RMB)	50.25	25.41	0.010	500.00
Demand	Total Number of Transaction	883.7	1.6E4	1	2.0E6

As we stated, one of the unique aspects of information that our data capture is the effect of the "conform or be cast out" policy. The platform announced the deadline for "conforming" as Week 10. Our data indicate that this policy was strictly implemented and significantly remodeled the market structure. In particular, by our calculating the number of sellers and the percentage of registered sellers, our data show that the platform started to enforce the expulsion policy beginning in Week 10 and took four weeks to expel sellers who had not registered. Figure 3.1 shows more than 20,000 sellers before the implementation of the policy but a sudden reduction of more than 10,000 sellers after implementation. One trend is that the number of registered sellers, although fluctuating, increases after the policy. The chart on the right shows that the platform starts to enforce expulsion beginning in Week 10 and continues to expel unqualified sellers in the following four weeks until the percentage reaches 100% in Week 14. Note that percentages after Week 14 are not exactly but very close to 100%. This is because newly entered sellers need an additional week(s) to go through the procedural registration process. The dots in the earlier time periods in the chart on the right show almost no changes of membership of platform refund insurance before Week 10, in support of the absence of forward-looking behavior with regard to registering the program before Week 10. Together with a pattern of an increase in the number of sellers after Week 10, the data imply that sellers might have no information about the potential policy change in earlier weeks and, thus, when the policy is implemented, have time only to strategically react after being expelled. The pattern of the increasing number of sellers also indicates that sellers might need to take time to react to the policy, implying a loss of equilibrium in the earlier periods after the policy shock.

Figure 3.1. Direct Impact of the "Conform or Be Cast Out" Policy

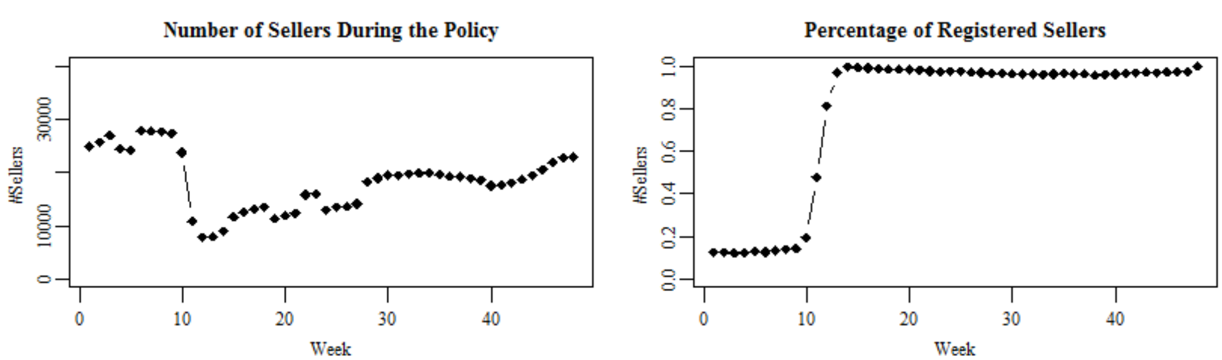

To understand the impact of the policy change, we further investigate the demand change during these periods. The chart to the left in Figure 3.2 shows platform-wide weekly demand. Here we find a sudden decrement of sales from Week 10, which lasts for four weeks. By comparing this chart with the chart on the left in Figure 3.1, we find a positive correlation between the number of sellers and total demand, more pronounced during Weeks 19 to 21 and Weeks 24 to 27, when there was a moderate fall in the number of sellers. This is consistent with our conjecture that variability positively affects platform-wide demand. Nevertheless, with regard to the magnitude of demand change when the policy was implemented, the charts present the decrement in demand around Week 10 as not being significantly larger than other demand fluctuations, especially when demand is compared to the change in the number of sellers. One explanation might be the selection effect of the policy, as expelled sellers tend to be small sellers who have limited impact on platform-wide demand. Another explanation might be strong substitution by conforming sellers who have improved quality based on the policy. Even though consumers might have less willingness to purchase on the platform due to decreased variety, the loss might be compensated through consumers' finding alternative sellers with improved quality with respect to platform refund insurance. If this explanation stands, we would expect an increment in demand for sellers who remain in the market and an even more pronounced increment for sellers who newly switched to the platform refund policy, as a substitution of expelled sellers occurs among all remaining sellers. This redistribution favors newly switched sellers due to improved quality, though early adopters are favored less due to escalated competition. Figure 3.2 also presents the average demand for a subset of sellers who are early adopters of the platform refund insurance before Week 10 and those who newly switch to the program in Week 10. As we expect, the demand after Week 10 for early adopters increase but not as significantly as that for newly switched sellers, shown as experiencing a peak from Weeks 10 to 15.

An exogenous policy change should result in sellers' strategic reactions until a new equilibrium is established. Figure 3.2 also presents increments of demand when approaching the new equilibrium. Interestingly, we find that the increment of demand is concurrent only with the market share change of early-adopter sellers. In contrast, demand for newly switched sellers, although surging when the policy is introduced, diminishes afterward until it reaches the level that it was before the introduction of the policy. Note that the average market share of early adopters is as large as 10 times of that of newly switched sellers. This suggests that, in the process of establishing a new equilibrium, early adopters who are also larger sellers regained their market share from small but newly switched sellers. Though we do not examine the formation of competition, the redistribution of demand leads us to suspect indirect effects of the policy, stemming from competition after the policy shock.

Figure 3.2. Demand Changes

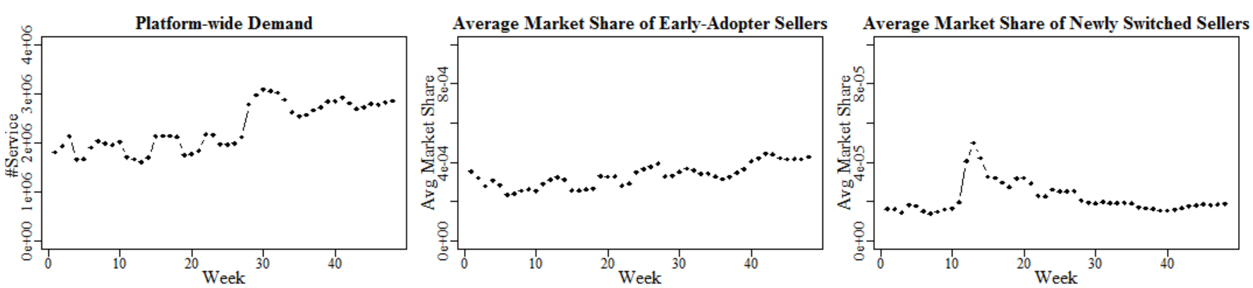

To be consistent with the demand estimation literature, the last item that we examine in a descriptive way is price. We present the log(price) distribution of the whole market, that of early-adopter sellers and that of newly switched sellers, respectively, in box plots with outliers in Figure 3.3. The charts show a slight shrinkage of price distribution when sellers are expelled but no apparent platform-wide mean level price change during implementation of the policy or afterward. With regard to sellers who choose to stay, we find a pattern of mildly increasing prices of early adopter sellers and newly switched sellers, suggesting a price markup for sellers who implement the entire policy. The markup for newly switched sellers is expected, as the platform refund insurance renders additional cost to newly switched sellers, whereas for sellers who adopted early, such markup indicates the potential indirect effect of the policy, e.g., escalated competition, which we address in a later section. Taken together, the data suggest the absence of price war, with exogenous policy change occurring.

Figure 3.3. Price Distribution per Week for All Sellers, Early-Adopter Sellers and Newly Switched Sellers

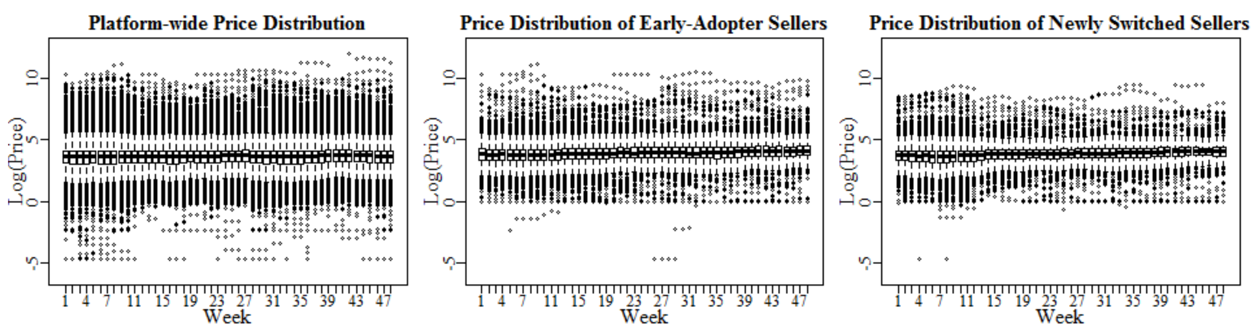

The descriptive findings provide us with insight into the demand system and the impact of the policy, but they are inadequate to quantify how and to what extent each characteristic in sellers' reputation systems and platform endorsements affects demand. The findings also fail to provide accurate estimates through which we can understand competition, consumer welfare, and the effect of the policy for the following reasons. First, there are a large number of sellers, and sellers have a very large action space. The change in demand might not be due solely to the change in one action but, rather, to a combination of many actions of other sellers. Without controlling other actions and states variables, we are not able to determine the exact effects of even one possible action if we observe a correlation. Second, sellers would take endogenous action toward the policy change. As a result, the demand change is a mix of the effect of exogenous policy change and that of endogenous actions. For a better understanding of the "conform or be cast out" policy, we need to distinguish the effects of exogenous change from those of endogenous actions as a means to measure the direct results of the policy and the indirect ones that occur after implementation. To this end, we build a BLP-style demand estimation model to recover the data-generation process and, thus, to provide accurate estimates that further help us to understand the "conform or be cast out" policy by counterfactual analysis.

3.4 MODEL

We present a model for the data-generation process for the market share of each seller in the marketplace. Specifically, we follow McFadden (1973) and, more directly, Berry et al. (1995) to account for unobserved consumer heterogeneity with a multinomial logistic model. We refer readers to Nevo (2000) for a detailed discussion of the methodological advantages of this model.

Suppose we observe sales data of each seller $j = 1,2, \quad ,J^t$ in each submarket defined as each time period $m = 1,2, \quad ,M$. A transaction is defined as the service of recharging a prepaid phone, which is differentiated with regard to service by seller in terms of delivery time and fulfillment rate. Consumers $i = 1,2, \quad ,I$ cannot have perfect information about quality of service; therefore, they rely on information provided by the platform to infer expected quality to help choose a seller. The conditional indirect utility that consumer i purchases service from seller j at market m is assumed to be of the form:

$$u_{ijm} = \beta_i X_{jm} + \xi_{jm} + \varepsilon_{ij}. \quad (3.28)$$

The first component in our utility function is X_{jm}, a vector that captures the multi-dimensional, observable characteristics of seller j. Assuming that consumers would access information for all potential sellers when they are seeking a service, consumers would receive signals from X_{jm} to form expectations of service quality of distinct sellers in a given market. There are three categories of information: past consumer-generated reputation, platform endorsement, and other seller-specific information. We include all numerically measurable variables relevant to sellers' reputation generated by past consumers as follows: (a) x_{1jm} = cumulative average reported delivery time; (b) x_{2jm} = log of cumulative count of disputes; (c) x_{3jm} = log of cumulative count of required refund; (d) x_{4jm} = merchandise average rating generated by consumers, with a scale of 1 to 5; (e) x_{5jm} = log of cumulative volume of ratings; (f) x_{6jm} = log of cumulative count of refunded transactions; (g) x_{7jm} = log of cumulative number of past transactions with positive evaluations; and (h) x_{8jm} = log of cumulative number of past transactions with negative evaluations. We include all possible combinations of platform endorsement for each seller as follows: (i) x_{9jm} = dummy for accepting a coupon tag; (j) x_{10jm} = dummy for a short-term VIP seller tag; (k) x_{11jm} = dummy for a threshold discount tag; (l) x_{12jm} = dummy for sample gift tags; (m) x_{13jm} = dummy for a seasonal short-term sale discount; (n) x_{14jm} = dummy for a VIP store tag; (o) x_{15jm} = dummy for accepting credit cards; (p) x_{16jm} = dummy for a detailed picture of the product; (q) x_{17jm} = dummy for platform refund insurance for potential consumers; and (r) x_{18jm} = dummy for guaranteed return and exchange within one week. The last type of information is determined purely by sellers. It consists of (s) x_{19jm} = log of the number of distinctive types of product, which are used to capture the variety and scope of products for a specific seller; (t) x_{20jm} = log of the number of days since becoming registered on the platform, which captures that duration of membership on the platform; and (u) x_{21jm} = average price of the product.

Other than the different characteristics of sellers, we incorporate further consumer heterogeneity by introducing heterogeneous tastes for different characteristics. We model taste parameters β_i as random coefficients to allow different individuals have distinct preferences toward a specific characteristic. Specifically, we follow BLP to model β_i, following a multivariate normal distribution:

$$\beta_i \quad \beta + \sum \upsilon_i, \tag{3.29}$$

$$\upsilon_i \quad MVN(0, I). \tag{3.30}$$

β is a 21 x 1 vector of mean parameters that are common across different individuals, which captures the mean sensitivities to different characteristics, whereas $\sum \upsilon_i$ captures individual level heterogeneity in their preference for characteristics. Following the BLP method, we model υ_i, following a multivariate normal distribution with a mean of zero and standard diagonal variance. Given that individuals might have different levels of variability toward different characteristics, we use $\sum$ to rescale the variance of taste coefficients.

Consumers might capture additional information of sellers that is unobserved by econometricians, e.g., offline reputation. We therefore denote ξ_{jm} as information about seller j in market m observed by consumers but not by us. This value is market and product specific and is common to all individual consumers. Consumers prefer a product with higher ξ_{jm} because it provides more expected utility.

Finally, consumers also have individual- and choice-specific unobserved information that is different across different consumers. We represent it as ε_{ij} and assume it to be independent and identically distributed across both products and consumers following a type-I extreme distribution. After rearranging the utility function in a more hierarchical manner and defining seller level mean utility as $\delta_{jm} = \beta X_{jm} + \xi_{jm}$ we have a suppressed format of utility:

$$u_{ijm} = \delta_{jm} + \sum \upsilon_i X_{jm} + \varepsilon_{ij}. \tag{3.31}$$

We accomplish the specification of the demand system by introducing utility of outside goods. Individuals might decide not to purchase the service from our online marketplace; instead, they might purchase a similar service from another online business or offline channel. We model the conditional indirect utility from the outside goods as:

$$u_{i0m} = c_0 + \upsilon_{i0}\sigma_0 + \varepsilon_{i0}, \tag{3.32}$$

where c_0 captures mean utility of outside goods, and $\upsilon_{i0}\sigma_0$ is an individual-specific component that reflects heterogeneity in the preference to choose outside goods. We model υ_{i0} as being drawn from a standard normal distribution $N(0,1)$ independently and σ_0 as a scaler parameter to be estimated. Another representation is to take c_0 as one more dimension of β with a constant covariate and take υ_{i0} as an additional dimension of υ_i, which leads the mean coefficient β to be a 22 x 1 vector, the taste heterogeneity distribution to follow 22 variates $MVN(0,I)$, and σ_0 to be an additional dimension of Σ. ε_{i0} is a "love of variety" individual-level, random error term that follows a type-I extreme distribution.

Assuming rationality of consumers, consumer i will choose seller j that generates maximal utility among all of the sellers in market m. According to our model above, the type of each individual can be characterized as $t_i = (\upsilon_i', \varepsilon_{i0}, \varepsilon_{ij})$ for $j = 1, 2, \quad, J$. The set of consumers in market m who choose product j can then be represented as:

$$A_{jm}(X_{\cdot m}, \xi_{\cdot m} \mid \beta, \Sigma) = \left\{ t_i \mid u_{ijm} > u_{ijt}, \forall j, t \;\; \text{s.t} \, j \neq t \right\}. \tag{3.33}$$

Given the property of type-I extreme distribution in a conditional multinomial discrete choice model, we would have a closed-form probability that consumer i would purchase product j in market m, according to:

$$\Pr_m\left(j \mid X,i\right)=\frac{\exp\left(\delta_{jm}+\sum\upsilon_i X_{jm}\right)}{\exp\left(c_0+\upsilon_{i0}\sigma_0\right)+\sum_{j=1}^{J}\exp\left(\delta_{jm}+\sum\upsilon_i X_{jm}\right)}. \quad (3.34)$$

Market share is obtained by aggregating over the population of potential consumers. Given that υ_i is distributed as 22 variates $MVN(0,I)$, we use $P(\upsilon_i)$ to denote a population distribution function of individual heterogeneity to taste coefficients. By taking integration over υ_i, we can further derive market share $s_{jm}\left(\delta_{jm},\theta\right)$ of seller j in market m as follows:

$$s_{jm}\left(\delta_{jm},\theta\right)=\int\Pr_m\left(j \mid X,i\right)d\left(P(\upsilon_i)\right)=\int\frac{\exp\left(\delta_{jm}+\sum\upsilon_i X_{jm}\right)}{\exp\left(c_0+\upsilon_{i0}\sigma_0\right)+\sum_{j=1}^{J}\exp\left(\delta_{jm}+\sum\upsilon_i X_{jm}\right)}d\left(P(\upsilon_i)\right), \quad (3.35)$$

where $\theta=\left\{\beta,\Sigma\right\}$.

3.5 ESTIMATION

3.5.1 *Identification*

Our approach resembles the generalized method of moment (GMM) approach proposed by Berry et al. (1991). The GMM-based BLP method essentially estimates the parameter to rationalize two sets of moment conditions. The first one equates market share predictions to the observed market share from our data, which is shown as follows:

$$s_{jm}\left(\delta_{jm},\theta\right)-s_{jm}=0 \text{ for } j=1,2,\quad ,J \text{ and } m=1,2,\quad ,M. \quad (3.36)$$

Berry (1994) shows the existence and uniqueness of δ_{jm} that guarantee this moment, under mild regularity conditions, on the distribution of consumer tastes.

The second moment condition concerns market-level disturbance. It requires unobserved market-level disturbance ξ_{jm} to be uncorrelated with observed exogenous variables and instrument variables denoted as Z and shown as:

$$E(\xi_{jm} \mid Z) = 0 \text{ for } j = 1,2, \quad ,J \text{ and } m = 1,2, \quad ,M . \tag{3.37}$$

The first moment guarantees a mapping between Σ and δ_{jm}. Given one specific δ_{jm} and the second-moment condition above, we could take β as a coefficient in a linear regression with δ_{jm} as a dependent variable and ξ_{jm} as random shock. This analogy indicates the identifiability of β, given δ_{jm}. With both conditions combined, we would have unique optimal extremum point identifying both Σ and β.

3.5.2 *Instrument Variables*

Similar to many empirical works on the online marketplace, our model is faced with the critical challenge of endogeneity. Endogeneity biases arise when firms are allowed to choose or adjust product characteristic X, given other information that is unobserved to econometricians ξ_{jm}. The traditional demand estimation literature (Berry et al. 1995; Petrin 2002) assumes price as the only adjustable strategic action by sellers in the short run, whereas it takes other characteristics as exogenous. Nevertheless, in the case of the online service marketplace, a much larger action space is adjusted on a relatively high frequency basis, which is unable to instrument itself. In fact, only past consumer-generated reputation and sellers' duration of membership are exempt from the simultaneity issue, as new consumers would have access only to the reputation system updated to time $m-1$, which is earlier than the arrival of unobserved demand shock at m. Sellers can adjust their platform endorsement characteristics simply by registering a corresponding service from the platform manager, implying that platform endorsement characteristics could be outcomes of strategic actions, given market and product demand shock ξ_{jm}. A similar argument could be applied to the price and variety of products, which are classic assumptions for a supply-and-demand system. In a more econometric interpretation,

$\left\{x_{9jm}, x_{10jm}, x_{11jm}, x_{12jm}, x_{13jm}, x_{14jm}, x_{15jm}, x_{16jm}, x_{17jm}, x_{18jm}, x_{19jm}, x_{21jm}\right\}$ are highly likely to be correlated with ξ_{jm}, which leads to harmful biasness toward our estimates.

We follow BLP and Berry (1994) to apply measures of isolation in product space as instrument variables. The idea is similar to cross-validation, which measures how different a product characteristic is from the market average characteristics, shown in the following:

$$z_{\cdot jm} = \frac{1}{J-1} \sum_{j' \neq j} x_{\cdot j'm} . \tag{3.38}$$

The BLP model shows that it is appropriate to use these instruments when price is the only endogenous variable and other characteristics could be taken as fixed and exogenous, which is less demanding than our setting, in which other characteristics also might be correlated with seller-specific shock. A similar justification, however, can be easily extended to our context. Note that our model focuses on a highly fragmented market, different from BLP models that focus on duopoly or oligopoly cases. This focus implies that each seller's unobserved characteristics have little potential to affect or reshape market average characteristics. In other words, even though seller-specific unobserved characteristics are likely to be correlated with seller-specific observed characteristics or those of a competing seller, they are unlikely to be correlated with market average characteristics. Therefore, we can still use a BLP-style instrument in our estimation.

One drawback of a BLP-style instrument is limited variation for certain variables. Thus, we supplement it with a second set of variables—Villas-Boas-Winer-style instrument variables (Villas-Boas and Winner 1999)—which uses lagged characteristics. Recall that we define each market as a market in a certain time period. The Villas-Boas-Winer-style independent variables in our example are, indeed, a special form of Hausman-style instrument variables (Hausman 1997) if we take the characteristics of a specific seller in another time as a special case of the characteristics of a similar seller in other markets. The intuition behind this is that a demand shock in time t might result in strategic adjustment of characteristics in time t; however, it will not result in an adjustment of characteristics in time $t-1$ due to the reversed order. Given that characteristics of a seller across time periods are correlated due to the continuity of strategic behavior and similar cost structure, lagged characteristics are valid instruments.

One potential issue with Villas-Boas-Winer-style instrument variables is seller-specific autoregressive demand shock over time. This results in a correlation between seller-specific demand shock in time t and that in $t-1$, which further leads to a correlation between current period demand and last-period price. Such an issue is alleviated in our model, however, as the persistent demand shock across time for a specific seller is already explained well by autoregressive exogenous variables, such as the cumulative number of positive evaluations and negative evaluations. Those variables ensure the elimination of a serially correlated component in ξ_{jm}.

Finally, we include lagged cumulative sales during time $m-1$ for each seller as an instrument. This information is visible to the consumer but is dropped from our main model due to very high collinearity with other characteristic variables that measure cumulative counts and the difficulty in normalizing other variables with regard to lagged cumulative sales due to a very high portion of sellers' having zero sales up to the last period. However, lagged cumulative sales is uncorrelated with seller-specific demand shock but highly correlated with the strategic action variables, including platform endorsement and pricing of sellers. Therefore, we include it as the third type of instrument variable.

3.5.3 *Estimation Specification*

Our model differentiates itself from other literature that applies BLP by the setting of fragmented markets. In fact, our data have almost 20,000 sellers, on average, in each market. This results in a much heavier computation burden, which requires us to be extremely cautious in implementing the estimation algorithm. There are two ways to estimate the BLP model. The original way, proposed by Berry et al. (1995), used contraction mapping to ensure the moment condition in equation 3.37 and applied GMM to estimate mean taste parameters. This algorithm requires a numerical solution in the inner loop to find δ_{jm} in each iteration, resulting in a heavy computational cost and potential numerical inaccuracy. Dubé et al. (2008) proposed to use an MPEC algorithm, which improves the computation performance by setting the inner loops as a constraint when solving the extreme point of the moment condition-derived objective function. The MPEC method, however, shows its weakness in a highly fragmented market when the number of players in each market is too large to result in a sufficient number of constraints in the numerical optimization (Dubé et al, 2008). In our example, the traditional BLP algorithm is at least 5 times faster than the MPEC method, which leads us to prefer the use of a contraction mapping algorithm in our study.

Specifically, with an initial value of Σ, we first solve for fixed points of δ_{jm}, which rationalize equation 3.36. Given the exponential converging rate of δ_{jm}, we would be able to solve this contraction mapping in a few minutes. Second, we apply the moment condition in equation 3.37 to solve for estimates of β , which minimize objective function conditional on δ_{jm}. We numerically iterate this step until we find an optimal point of Σ that generates δ_{jm} that minimizes overall objective function with estimates of β . The second unique side of our estimation specification concerns standard error. Note that the asymptotic variance of $(\hat{\beta}_{GMM}, \hat{\Sigma}_{GMM})$ results in computing the Jacobian of the moment conditions with respect to (β,Σ). For $\hat{\beta}_{GMM}$, the Jacobian simply reduces to X. However, for Σ, the estimation requires computing inversion of the following J by J matrix:

$$\left[\frac{\partial s_{jm}\left(\delta_{jm},\theta\right)}{\partial \delta_{j'm}}\right]_{j=1\quad J,\,j'=1\quad J}. \tag{3.39}$$

Note that when we have, on average, 18,306 sellers in each market, the inversion of a matrix as large as 18,306 x 18,306 is very computationally demanding. The calculation, although computationally doable, results in loss of information and, thus, inaccuracy due to pseudo-inversion. To handle this issue, we apply a bootstrap technique to recover the variance parameters. The bootstrap technique provides an advantage by allowing parallel computing. Even with 100 iterations of bootstraps, we need to spend only as little as 10 times the duration for each point estimate optimization when we have 10 parallel lines. The computational gain would be augmented when we have more lines.

3.6 Estimation Result

In this section, we present the estimation results and insights associated with the resulting numbers. As noted earlier, the BLP-style model allows researchers to analyze not only demand for sellers but also welfare for consumers and the overall impact on the platform. Therefore, we present estimation results in three parts: Section 3.6.1, parameter estimates and marginal effects; Section 3.6.2, substitution effects; and Section 3.6.3, consumer welfare analysis.

3.6.1 *Parameter Estimates and Marginal Effects*

In Table 3.5, we present mean and variability coefficients for each characteristic, along with standard errors. When interpreting the estimated model, researchers should note the good properties of a random coefficient logit model that allows distinct sellers to have different marginal effects for a specific variable. The random coefficient logit model incorporates more degrees of freedom of the model to fit in and, thus, leads to a more realistic interpretation of the results. Confusion results, however, when the sign of the mean parameters is not consistent with the sign of the marginal effect due to nonlinearity of the multinomial logit transformation and the integration of random coefficients. Note that linear models always have marginal effects consistent with mean effects because the integration of random coefficients would always be equal to mean effect. A nonlinear transformation of random coefficients would make the marginal effect deviate from that of the model in terms of the mean coefficient only. In other words, the variability parameters infuse an effect in addition to the one from mean parameters when a nonlinear transformation is imposed. When the sign of the variability effect is opposite to that of the mean effect, the marginal effect will be opposite to the mean effect if the variability effect shows more strength. Therefore, researchers need to be extremely cautious when interpreting a coefficient with large variability parameter.

We present the marginal effect of each variable in addition to the estimate of the mean parameter and that of the variability parameter to avoid any potentially misleading results. Because the marginal effect of each characteristic is seller and market specific, we calculate the marginal effect of several representative sellers, including a megaseller who has a market share of over 5%, a median seller, and a 25% quantile small seller in market $m = 1$ (we name them as representative sellers, interchangeably). The margin we use is a 0.1 increment for average rating and log(110%) increment for a log-transformed variable in Euclidean distance. Because dummy variables are not differentiable, we present marginal effects by doing a *what-if* analysis to show the differences between the effect of data-observed choice and that of an alternative choice. By exhibiting marginal effect, we provide a more direct and explicit understanding of the impact of each characteristic on market share, as compared with interpreting the value of coefficients only. We summarize the numerical results as a percentage of marginal demand relative to current demand for representative sellers, as seen in the last three columns of Table 3.5.

Table 3.5. Estimation Results

Variable	Base Coefficients		Variability Coefficients		Marginal Effects		
					Mega Seller	Median Seller	0.25 Quantile Seller
Delivery time	-1.21*	(0.97)	0.88*	(0.51)	-4.26%	-9.05%	-12.20%
#Dispute	-0.11***	(0.03)	0.27***	(0.08)	-0.32%	-1.30%	-1.84%
#Refund RQST	-0.11*	(0.09)	0.16***	(0.04)	-0.05%	-0.86%	-1.13%
Avg rating	0.50**	(0.22)	0.06***	(0.01)	1.86%	3.76%	2.28%
Rating volume	0.97***	(0.18)	0.11***	(0.03)	3.72%	6.98%	9.64%
#Refund	-0.19***	(0.04)	0.56***	(0.07)	2.28%	-2.36%	-3.36%
#Positive EVAL	0.69***	(0.23)	0.03*	(0.02)	2.64%	4.99%	6.83%
#Neg EVAL	-0.09***	(0.03)	0.08***	(0.02)	-0.32%	-0.58%	-0.73%
Coupon	0.02	(56.36)	1.06	(25.47)	-5.16%	70.47%	96.79%
Short-term VIP	-0.33**	(0.17)	1.01***	(0.28)	2.17%	-4.46%	-6.15%
TD	-0.11	(0.14)	1.03***	(0.29)	11.66%	37.36%	52.90%
Sample gifts	-0.70***	(0.20)	1.02***	(0.26)	-22.39%	16.97%	25.40%
Short term sale	-1.60***	(0.40)	0.85***	(0.14)	-55.04%	-52.07%	-71.35%
VIP store	0.23*	(0.18)	0.86***	(0.21)	8.81%	56.70%	84.79%
Credit card	-0.17*	(0.10)	0.95***	(0.22)	8.24%	11.46%	14.06%
Detailed picture	0.22	(0.28)	1.03**	(0.42)	20.39%	123.08%	166.29%
PRI	0.64***	(0.15)	0.23***	(0.06)	21.12%	75.60%	103.43%
SGR	-0.44***	(0.08)	0.93***	(0.19)	-35.80%	9.45%	10.93%
Product variety	0.12**	(0.05)	0.01***	(0.00)	0.44%	0.82%	1.12%
Duration	0.40**	(0.13)	0.10*	(0.06)	1.55%	2.86%	3.98%
Price	-3.04**	(0.99)	0.01**	(0.00)	-10.84%	-18.41%	-25.17%
Outside goods	2.35	(3.46)	0.03*	(0.03)	-	-	-

*Z statistics > 1; ** Z statistics > 2; ***Z statistics > 3. Standard error in parenthesis.

Our results show the effectiveness of consumer-generated reputation. Most coefficients are significant with the expected signs. All else equal, consumers are qualitatively more inclined to consume from sellers with faster service speed, fewer disputes, fewer required refunds and fulfilled refunds, higher average ratings, larger rating volume, and more positive evaluations and fewer negative ones.

Most of the marginal effects of consumer-generated reputation are consistent with the sign of the mean parameters, given that variabilities are mostly moderate to small, with the exception of cumulative fulfilled refund count (log), whose variability parameter (0.56) is at least three times larger than the mean parameter (-0.19), implying very fat tails of coefficient distribution and, thus, very heterogeneous tastes for this variable. One possible explanation for the high variability is that refund count, on one hand, has a negative impact on consumers' utility by representing the risk of extra cost associated with a failed transaction; on the other hand, it signals its positive side as it implies sellers' willingness to undertake responsibility. The variability leads to contrasting marginal effects of a fulfilled refund for different sellers, with -2.36% and -3.36% for median and small sellers but 2.28% for megasellers, with the sign flipping when cumulative refund counts increase by 10%. This suggests a marginal positive effect from the mass of consumers who favor refund counts that dominate the marginal negative one for a big seller.

For other consumer-generated reputation, the signs of marginal effect are consistent with intuitions and mean parameters, and the scale of marginal effect relative to size of the seller is more pronounced for smaller sellers, whereas the exact marginal demand changes are higher for larger sellers. In particular, among different characteristics, reported delivery time is among the highest marginal effect, as slowing down the speed 1.1 times would result in a 4.26%, 9.05%, and 12.20% demand decrease for mega, median, and small sellers, respectively, followed by rating volume, as a 3.72%, 6.98%, and 9.64% demand increase for a 10% additional rating. Lowering dispute counts by 10% will increase demand by 0.32%, 1.30%, and 1.84%, respectively, for the three sellers. Similarly, a 10% reduction of the requirement for refund draws a 0.05%, 0.86%, and

1.13% demand for representative sellers. Increasing the average rating by 0.1 will attract demand significantly, by 1.86%, 3.76%, and 2.28% for representative sellers. Finally, a 10% accumulation of positive evaluations increases next time-period sales by 2.64%, 4.99%, and 6.83 for representative sellers, whereas that for negative evaluations lowers demand by 0.32%, 0.58%, and 0.73%.

Platform endorsement, in contrast, shows very heterogeneous tastes of consumers, except in terms of platform refund insurance, in support of our conjecture that rational individuals would be able to perceive the cost of endorsements as a markup in the price, and the benefits of endorsement would be applied to only a subset of targeted consumers. Recalling that the sign of mean parameters might be inconsistent with that of marginal effects when variability is large, the interpretation of mean parameters might not be as meaningful as the marginal effect.

Our calculations suggest that the effect of variability parameters for platform endorsement could be large when contrasting the sign of marginal effect with that of mean parameters. Specifically, the endorsement with lower variability parameters have a sign of marginal effect consistent with that of mean parameters, e.g., having a VIP store results in 8.81%, 56.70%, and 84.79% additional demand for representative sellers; having a detailed picture of product increases demand by 20.39%, 123.08%, and 166.29% for representative sellers; platform refund insurance also positively affects demand by 21.12%, 75.60%, and 103.43% for megasellers, median sellers, and small sellers, respectively. Surprisingly, a short-term sales discount tag is shown as negative for demands with -55.04%, -52.07%, and -71.35% marginal effects for representative sellers and negative mean parameters. The reason might be that a sales tag is no longer attractive, given that price is transparent. In contrast, platform endorsement with higher variability parameters exhibits marginal effects with the opposite sign to the mean parameter for some or all of the sellers. In particular, a threshold discount tag, although shown as negative in the mean parameters, increases demand by 11.66%, 37.36%, and 52.90% for representative sellers. Similarly, a tag that indicates acceptance of credit cards generates 8.24%, 11.46%, and 14.06% additional demand, in contrast to its sign of the mean parameters.

Our results also show monotony of marginal effects with regard to size of sellers, indicating the possibility that only one end of sellers in the size dimension exhibits a contrasting marginal effect. Specifically, we find that a coupon tag, although decreasing demand by 5.16% for megasellers, increases demand by 70.47% and 96.79%, respectively, for median sellers and small sellers; the short-term VIP seller tag, although increasing sales of megasellers by 2.17%, decreases that of median and small sellers by 4.46% and 6.15%, respectively; a sample gift tag shows the positive effect on demand by 16.97% for median sellers and 25.40% for small sellers but lowers demand by 22.39% for megasellers. Seller-guaranteed return and exchange tags foster demand for median and small sellers by 9.45% and 10.93%, respectively, whereas they decrease megasellers' demand by 35.80%, correspondingly. These findings suggest that platform endorsement for small and median sellers might be harmful to megasellers and vice versa. They also indicate the potentially incorrect qualitative interpretation of estimation results if researchers focus only on the sign of the mean parameters, addressing the importance of incorporating random coefficients in demand estimation.

We find expected effects of other seller-specific characteristics. With a moderate variability parameter estimate, the sign of mean parameters exhibits consistency with that of marginal effects. Product variety positively increases demand by 0.44%, 0.82%, and 1.12% for representative sellers. A 10% longer duration of membership also increases sales by 1.55%, 2.86%, and 3.98% for representative sellers, implying consumers' preference of sellers with a longer transactional history. As expected, a 10% increase in price lowers demand by 10.84%, 18.41%, and 25.17% for mega, median, and small sellers, respectively.

3.6.2 *Substitution Effects*

A strategic characteristic adjustment by sellers would not only have an impact on the seller but also would exhibit externalities by affecting demand for competitors. To understand the externalities in the online service marketplace, we calculate self- and cross-elasticity of demand for continuous variables. Given the large number of sellers and high dimensional characteristics, we apply the calculation only on reported delivery time and price for selected representative sellers as 0.25 quantile (small) seller, median seller, 0.75 quantile (large) seller, and megaseller in market $m=1$ according to:

$$\eta_{\cdot jkm} = \frac{\partial s_{jm} x_{\cdot km}}{\partial x_{\cdot km} s_{jm}} = \begin{cases} -\frac{x_{\cdot km}}{s_{jm}} \int \beta_{\cdot i} s_{ijm} \left(1 - s_{ijm}\right) d\left(P(\upsilon_i)\right), & \text{if } j = k; \\ \frac{x_{\cdot km}}{s_{jm}} \int \beta_{\cdot i} s_{ijm} s_{ikm} d\left(P(\upsilon_i)\right), & \text{otherwise,} \end{cases} \tag{3.40}$$

where $\eta_{\cdot jkm}$ measures the responsiveness of the demand for seller j with a change of characteristics of seller k. s_{jm} is the market share of seller j in market m. $x_{\cdot km}$ is the changing variable of seller k in the same market. $\beta_{\cdot i}$ is a consumer-specific parameter associated with the changing variable. s_{ijm} (s_{ikm}) are succinct representations of individual-specific willingness to purchase as $\Pr_m\left(j \mid X, i\right)$ ($\Pr_m\left(k \mid X, i\right)$), as shown in equation 3.34. For an explicit interpretation, we additionally calculate counterfactual marginal effects of representative sellers to measure the changes in market share if one seller were to make a strategic adjustment. In Table 3.6 and Table 3.7, we take the characteristics of the seller in the head of each row as changing variables and calculate marginal effects and elasticity to measure the responsiveness of the demand for the seller in the head of each column.

Delivery time was investigated first due to its strongest marginal impact among consumer-generated reputation. Coupled with negative self-elasticity, positive cross-elasticity implies that decreasing reported delivery time would allow the seller to substitute demand squeezed from the market share of other sellers. When lowering reported delivery time by 10%, the demand for the seller increases proportionally to the self-elasticity and size of the seller, whereas the demand for the competitor drops proportionally to its own size and the increment of the changing sellers' demand. Therefore, the larger the size of a changing seller, the stronger the externalities that it will exhibit; the larger the size of the responding seller, the stronger the impact that it will receive.

As a classic issue in economics, we measure price elasticity of demand and the changes in demand when a changing seller drops its price. Consistent with the properties of a competitive market, we find large and positive price self-elasticity of demand and negative cross-elasticity, suggesting a price war as an effective strategy to squeeze demand from competitors. By examining the scale of market share changes when the changing seller provides a 10% price discount, our results show that, compared with smaller sellers, the larger the size of the seller, the more additional market share it will gain, resulting in a greater loss of competitors when it is the changing seller, and the stronger negative impact it will receive if the competitor is the changing one.

Table 3.6. Delivery Time (log) Substitution Pattern

Delivery Time	Marginal Market Share with 10% Decrease in Delivery Time				Delivery Time Elasticity of Market Share			
	Mega	Large	Median	Small	Mega	Large	Median	Small
Mega	0.00185	-4.94E-11	-1.82E-11	-7.56E-12	-0.1328	3.41E-09	1.26E-09	5.17E-10
Large	-5.04E-11	8.74E-07	-3.00E-12	-1.05E-12	2.32E-05	-0.4113	1.41E-06	4.74E-07
Median	-1.86E-11	-3.00E-12	3.08E-07	-3.34E-13	2.43E-05	1.48E-06	-0.416	4.31E-07
Small	-7.77E-12	-1.05E-12	-3.34E-13	1.11E-07	3.01E-05	4.06E-06	1.30E-06	-0.416

Table 3.7. Price Substitution Pattern

Delivery Time	Marginal Market Share with 10% Decrease in Price				Price Elasticity of Market Share			
	Mega	Large	Median	Small	Mega	Large	Median	Small
Mega	0.004914	-1.68E-10	-5.86E-11	-1.72E-11	4.300275	-1.52E-07	-4.77E-08	-1.57E-08
Large	-1.72E-10	1.98E-06	-5.93E-12	-2.08E-12	-9.45E-04	12.02692	-3.25E-05	-1.24E-05
Median	-6.00E-11	-5.93E-12	6.94E-07	-6.69E-13	-9.40E-04	-1.03E-04	10.87022	-1.14E-05
Small	-1.81E-11	-2.08E-12	-6.69E-13	2.41E-07	-0.00082	-0.0001	-3.02E-05	11.87648

3.6.3 *Consumer Welfare Analysis*

Our model allows us to analyze the impact of changes in the market structure on the consumer side by providing a closed-form solution for consumer surplus. Therefore, we calculate the individual consumer surplus across different time periods to examine how a compound market structure change affects consumers. Following McFadden (1981), we calculate the average individual-level consumer surplus in a random coefficient logit model:

$$CW_m = \int \frac{\ln\left(\sum_{j=0}^{J_m} \exp\left(u_{ijm}\right)\right)}{-\beta_{21i}} d\left(P(\upsilon_i)\right), \tag{3.41}$$

where u_{ijm} is the individual-specific indirect utility function shown in equations 3.31 and 3.32, J_m is the number of sellers available in market m, and β_{21i} is the individual-specific coefficient for price, which is equal to mean parameter β_{21} plus individual-specific deviations from the mean parameter $\Sigma_{21}\upsilon_i$. By dividing indirect utility by marginal effects of income β_{21i}, CW_m estimates the consumer welfare in a monetary unit. Taking the difference between CW_m in market m and CW_0 as consumer welfare in a benchmark market structure would lead to compensating variation (CV) and equivalent variation (EV). Therefore, comparing CW_m allows us to understand whether consumers are better off, given a specific market structure.

We present the consumer welfare along with time (markets) in Figure 3.4, which shows that an average consumer would gain 0.5 to 2 (RMB) from the availability of the online service marketplace. As time passes, consumers are shown to be better off, with a trend of increasing consumer surplus. Recall that the platform started the “conform or be cast out” policy at Week 10, and our result is consistent with past literature that shows increased consumer welfare with higher levels of product variety (Brynjolfsson et al. 2003, Ghose et al. 2006) in a pattern of decreasing consumer surplus when sellers are expelled after Week 10. Nevertheless, consumer surplus rebounds as the market structure moves to a new equilibrium after the exogenous shock.

Figure 3.4. Average Consumer Welfare

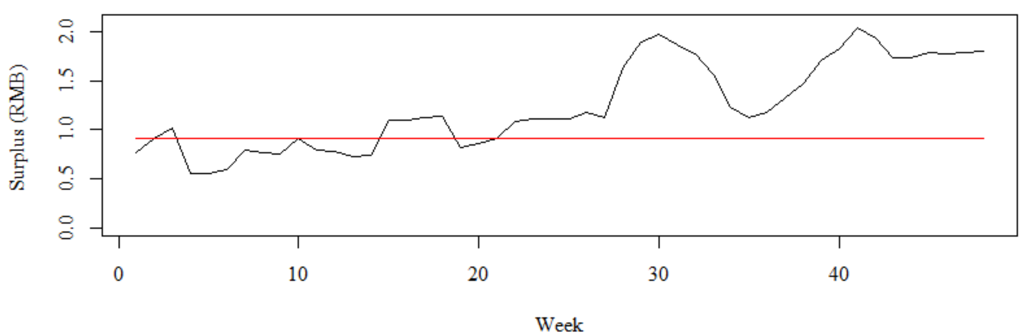

Note that the total number of sellers rebounds as well after policy implementation but is still much less than before implementation. Researchers might draw the conclusion that platform refund insurance plays a more pronounced role than does seller variety with respect to consumer welfare. However, such a conclusion is hasty, as it potentially overvalues the direct effect of policy change and overlooks the indirect effect of the endogenous reactions of sellers. One possible indirect effect is the escalation of competition. One typical way to analyze sellers’ strategic behavior after a policy shock is to derive a theoretical optimal action, assuming that all sellers are profit maximizers. This approach contradicts our observation, however, as sellers are not rational enough to behave optimally immediately; further, this approach is theoretically impossible to solve with a high dimension of action space and a larger number of players in our context. We therefore take an alternative approach, using counterfactual simulation to approximately decompose the overall effect of the policy on consumer welfare into direct effects and potential indirect effects to deepen our understanding of the “conform or be cast out” policy.

3.7 Policy Simulation

We conduct three sets of policy simulations with the techniques that we developed in former sections to understand the impact the “conform or be cast out” policy. Recall that the policy requires sellers to either register for platform refund insurance or exit the market. The policy makes a two-sided direct impact on the demand for sellers who stay, which includes: (1) externalities due to improved service quality of sellers who switch from a state of not being insured to that of being insured, and (2) substituted demand from reduced seller variety. In addition, strategic sellers might adjust their characteristics, including platform endorsement, price, and product variety, after the policy change, which is an indirect effect of the policy. Therefore, observing even an overall increment of demand after the policy, we still await the attribution of different effects.

Specifically, we note that most markets are highly likely to lose equilibrium due to their proximity to exogenous shock. Thus, we use the two tails of the observed market (Market 1 and Market 48) to set our experiment, as they are closer to equilibrium before and after the policy change. In the first policy simulation, we investigate the direct effect of platform refund insurance, and, in the second one, we investigate that of the expulsion rule. Given that a realistic policy includes a combination of expulsion and platform refund insurance registration, our last policy recovers a counterfactual market with the same direct effect of the policy as that for Market 48 but based on the original equilibrium for Market 1. We finally explain the indirect effect by comparing a simulated market with Market 48 that includes both a direct and indirect effect in the new equilibrium.

3.7.1 *Direct Effect of "Conform"*

We first investigate the marginal direct effect of platform refund insurance. In particular, we conduct a counterfactual simulation to investigate the impact on sellers and calculate consumer welfare to provide insight into the perspective of the platform and consumers. To understand the marginal self-effect and externalities, the counterfactual analysis simulates the market share of all sellers when one specific seller switches its state of platform refund insurance and calculates the marginal effect by differentiating market share between the original market and the simulated one, as shown in Table 3.8, where the head of each row is the changing seller. Due to limited space, we report effects on only representative sellers, including megasellers, 0.75-quantile (large) sellers, median sellers, and 0.25-quantile (small) sellers as well as the overall externalities denoted in the "All Others" column. Table 3.8 shows that the platform refund insurance, on the one hand, increases the demand for sellers who adopt it and, on the other hand, lowers the demand for other, unchanged sellers. With respect to the scale, the larger the size of the seller, the more additional demand it will gain and the stronger the negative externalities it will impose if adopting this endorsement by itself and the stronger the negative impact it will receive if the other seller(s) is the changing one.

Table 3.8. Direct Effect of Platform Refund Insurance

Platform Refund	Demand Marginal Effect on Seller					Consumer Welfare
Insurance	Mega	Large	Median	Small	All Others	CV(EV)
Mega	0.00858	-2.19E-10	-7.77E-11	-2.60E-11	-0.00856	0.0099
Large	-3.74E-10	6.09E-06	-1.84E-11	-5.12E-12	-2.26E-06	7.03E-07
Median	-1.34E-10	-1.84E-11	2.13E-06	-1.80E-12	-7.78E-07	7.03E-07
Small	-2.29E-11	-2.40E-12	-8.43E-13	2.29E-11	-1.04E-07	9.31E-08
	Already Registered		Newly Registered			
All	-2.68%		9.38%			0.0408

We further calculate consumer welfare by taking the difference in consumer surplus between the original setup and the counterfactual setup, shown in last column of Table 3.8. As expected, our result indicates that sellers' adoption of platform refund insurance will benefit consumers, and adoption by a larger seller generates higher levels of consumer welfare.

To understand the overall direct impact when all sellers are required to register for insurance, we further simulate a case with all sellers registered, as shown in the last two rows of Table 3.8. Our results indicate that, when all sellers are registered, sellers who had already been registered are worse off due to increased competition and that sellers who are newly registered are better off, given their improved competency. Overall, the policy creates 6.70% additional demand for the platform as a net of 2.68% loss of demand for sellers who have registered and a 9.38% increment from newly registered sellers. Given that the quality of every sellers' service is improved with platform refund insurance, consumers are supposed to be better off, which is validated by a 0.0408 increment of consumer surplus.

3.7.2 *Direct Effect of "Being Cast Out"*

The other direct effect of the policy is the "casting out" of unconfirmed sellers. We take a similar approach to examine the direct effects of the expulsion of sellers as that in Section 3.7.1. Specifically, we simulate four markets with the exit of a megaseller, exit of a 0.75-quantile seller, exit of a median seller, and exit of a 0.25-quantile seller and examine subsequent demand changes for the other sellers and overall externalities. In addition, we simulate a more realistic setup with exits of many sellers, equal to that of Market 48, to measure the overall impact of the exit of multiple sellers. In this way, we assume rationality of sellers such that action will be immediately taken, instead of there being a time lag, as we observed. The sellers who we manipulate to exit are randomly drawn from sellers who have, indeed, left during the time period. We measure the effect of "being cast out" by calculating the change in market share before and after the exit. Moreover, we calculate compensating variation to measure the change in the welfare of consumers. The results are shown in Table 3.9, with the head of each row as representing the expelled seller.

Table 3.9. Direct Effect of Seller Exit

Expelled Sellers	Demand Marginal Effect on Seller					Consumer Welfare
	Mega	Large	Median	Small	All Others	CV(EV)
Mega	-0.04062	5.62E-10	1.96E-10	7.03E-11	0.040567	-0.02775
Large	5.13E-10	-5.90E-06	1.76E-11	6.17E-12	2.19E-06	-1.94E-06
Median	1.79E-10	1.76E-11	-2.06E-06	1.99E-12	7.50E-07	-6.80E-07
Small	5.38E-11	6.17E-12	1.99E-12	-7.16E-07	2.81E-07	-2.36E-07
	Expelled		Stayed			
1944 Sellers	-3.22%		2.11%			-0.012

Our results indicate that losing the benefits of the expelled sellers, however, hurts the welfare of consumers as well as the overall market share of the platform. The larger the size of the exiting seller, the more marginal benefits that other sellers would gain and the more that consumers would lose. Note that Market 1 has 1,944 more sellers than has Market 48. To control the market structure, we mimic the policy by expelling 1,944 sellers who are randomly sampled from sellers who, indeed, exit. The results show 0.012 (RMB) losses of consumer surplus. Even though expulsion increases the overall market share of sellers who remain by 2.11%, the loss of the market share of expelled sellers leads to a net market share loss of 1.11% for the platform, which is consistent with our observation in Figure 3.2.

3.7.3 *Direct Effect of the "Conform or Be Cast Out" Policy and its Potential Indirect Effects*

In this section, we combine two aspects of the policy that we examined in Sections 3.7.2 and 3.7.3 to investigate the overall direct effect when sellers are not allowed to make any strategic adjustment and further derive indirect effects when sellers' reactions are considered. Specifically, we control the change of market structure due to the policy by expelling 1,944 sellers and requiring all sellers to register for platform refund insurance. As a result, the seller variety and platform refund insurance registration rate of the counterfactual market are the same as those of Market 48. Note that we assume that sellers are not allowed to strategically react to the policy change. Therefore, other characteristics of sellers in the simulated market are the same as those in the original Market 1, which is in the equilibrium before the policy change, thus allowing us to examine potential indirect effects by comparing the simulated market with the observed Market 48. Because the direct effects of the policy are well controlled, all the differences come from the other characteristics, which are not directly affected by the policy but are strategically adjusted by sellers, given the policy change.

We measure the impact of these direct effects on different stakeholders of the platform. Sellers who newly register for platform refund insurance are certainly better off with respect to demand because both aspects of the policy lead to increments in demand. Our simulation shows an overall 9.33% market share increment for this party. For sellers who were registered before the policy change, there are two contrasting effects of the policy, which are shown in Sections 3.7.1 and 3.7.2. The exact demand change will be dominated by the aspect of the policy that shows more strength. Our findings show controversial direct effects among individual sellers. A summation, however, shows a general 0.72% loss of market share for this party, indicating a worse-off condition for sellers who were registered for the program, as the additional market share gain from expelling sellers appears inadequate to offset the loss from the escalation of other sellers.

The direct effect makes the whole platform lose a market share of 3.22% by expelling sellers who refuse to register for platform refund insurance; however, an additional 8.61% market share is gained from the improved seller quality, which results in an overall 5.39% increment in market share. This is contrary to what we observed, as seen in Figure 3.2, because the simulated policy incurs only loss of variety from 1,944 sellers, which is only one-fifth of the loss in the observation. This implies that, without considering sellers' strategic adjustments, the simulated policy change makes the market better off.

We calculate consumer surplus for the counterfactual market and derive compensating variation through a comparison with the original Market 1. The results show CV as 0.027 (RMB), implying that improved quality sufficiently offset the loss from less variability, which leads to consumers' being overall better off.

Finally, given the measure of the direct effect of the policy, we could measure the indirect effect by comparing the simulated market and Market 48. Our results show that, in addition to the increment in market share of 5.39%, the indirect effect further increases the market share by 25.36%. With regard to consumer welfare, if we set Market 48 as the initial policy, CV is shown as -1.02 (RMB). Note that the difference between Market 48 and the simulated market is about only characteristics other than platform refund insurance and seller variety, which are assumed to be close to equilibrium before and after policy implementation. Our findings indicate that the platform and consumers are significantly better off in the new equilibrium after the policy, suggesting a potentially positive indirect effect of the policy if there is no other policy change or positive exogenous shock to the market that we do not observe. The negative CV suggests that the overall characteristics of sellers in the marketplace are improved, which leads to a higher level of utility that consumers could gain from consumption in the platform. Given that improvements in characteristics might result in increasing costs, this finding is consistent with the markup of price of early-adopter sellers, as shown in Figure 3.3.

Although our method is not adequate to provide a conclusive explanation for the positive indirect effect, we conjecture that the policy, in general, homogenizes sellers by expelling and implementing the same characteristics and triggers new competition for sellers who have the motivation and ability to differentiate themselves from each other, which finally escalates the overall characteristics in the new equilibrium. Given that competition results in further cost for sellers, we cannot draw a conclusion in regard to whether sellers are overall better off with respect to profit, without knowing how much such competition costs. For the platform and consumers, however, the indirect effects result in higher market share and consumer welfare, suggesting the success of the policy for platform managers with a goal of escalating the overall performance of the platform. Given that our paper focuses on demand-side analysis, we do not investigate the formation of the new equilibrium or the exact path of improvement of sellers' characteristics. We leave these issues for future studies.

3.8 Conclusion and Implications

Unlike most physical goods, services sold in the online marketplace come with limited information about quality such that consumers are highly dependent on the reputation system and platform endorsement to choose sellers. Research is needed to help sellers and platform managers to quantify the returns based on the reputation system and platform endorsement in service marketplaces and to examine the effectiveness of related policy with the goal to improve the overall quality of consumption on the platform.

We apply a BLP-style model that recovers the demand data-generation process to study how consumer-generated reputation and platform endorsement affect the purchasing behavior of consumers and, consequently, the demand for sellers. Applying the model to aggregate-level sales data for recharging in a prepaid-phone service, we find that consumer-generated reputation significantly motivates consumers to purchase from sellers with faster delivery rates, fewer disputes, fewer required refunds and fulfilled refunds, higher average ratings, larger rating volume, and more positive evaluations and fewer negative ones. In addition, sellers' sensitivities toward information from the reputation system are quite homogeneous, except toward the number of fulfilled refunds. In contrast, the distribution of consumers' tastes for platform endorsements, such as coupons, short-term VIP seller tags, threshold discounts, sample gifts, accepting credit cards, and guaranteed return and exchange, is quite widespread, resulting in inconsistent marginal effects on demand across different sellers. One possible explanation might be that the benefit of a certain type of platform endorsement would be enjoyed only by a subgroup of matching consumers who favor the endorsements, while the rest of consumers would perceive platform endorsement as transferred cost and as resulting in a negative impact. With regard to platform endorsement in the case of less heterogeneity in taste, demand increases consistently when sellers register tags for platform refund insurance, VIP stores, and detailed pictures and decreases with tags of seasonal short-term sale discounts. Further, we derive self- and cross-market share elasticities of sellers' characteristics to quantify the impact of competition when sellers adjust specific characteristics and consumer surplus to understand the overall performance of the platform in monetary units.

We use the estimate of our model to conduct an empirically oriented policy analysis of the "conform or be cast out" policy that is observed in our data. With a counterfactual simulation, we measure the direct effect on demand and consumer welfare of platform refund insurance and that of the expulsion rule. We find the externalities of both aspects of the policy, as sellers who switch to the insurance plan squeeze demand from the market share of sellers who are already in the plan, and sellers who decide to conform and stay cannibalize the market share from sellers who exit the market. Consumers would gain additional surplus with an improved overall insurance policy of sellers, while lose some welfare due to diminishing seller variety. By combining both aspects of the policy, we find that sellers who newly switch to insurance plans would gain additional market share. The marginal impact on sellers who are already in the plan, however, varies, depending on whether the cannibalized market share from expelled sellers is enough to offset the loss from the escalated competition of platform refund insurance. Consumers are better off, given that premiums from improved insurance adequately cover the loss from a smaller selection set. Further, by controlling the state of the platform refund insurance and the variety of sellers who remain in the market, we compare markets before and after policy implementation to infer potentially indirect effects that are generated by strategic adjustment of other characteristics by sellers after the policy shock.

Our results imply that the policy triggers further competition in which sellers upgrade their characteristics. The competition, in contrast, improves the overall quality of the platform and increases consumer welfare by about 1 (RMB). Given that the policy, in general, shortens the distance between sellers in characteristic space, one possible explanation for the escalation is that sellers strategically seek to differentiate themselves after the homogenizing shock. To our knowledge, this empirical study is the first to examine explicitly the effectiveness of the "conform or be cast out" policy in the online marketplace. Although these estimates are based on the assumption of stationary, unobserved shock, they nevertheless provide the best possible estimates of further reactions of sellers and represent the process of approaching a new equilibrium after the policy change.

There are some limitations of our research. First, a more accurate and efficient model could be developed if a supply side function were incorporated. Although it may not be feasible to model a fully two-sided market, it would allow us to endogenize platform endorsement variables and sellers' other decisions and further allow us to understand the strategic action of sellers as well as the formation of equilibrium explicitly. However, to do so poses a methodological challenge in our context, given the very large action space and very large number of players. The nonstationary evolving of the characteristics of sellers also indicates potential off-equilibrium decisions of sellers, which cast doubt on the validity of the profit maximizer assumption. We leave this issue to be addressed in empirical studies. In addition, our model has limitations with regard to the entry and exit decisions of sellers. Future research can incorporate these decisions to provide a more accurate measure of the effect of "casting out" and the overall platform evolution by endogenizing these decisions. Further, future research should investigate individual-specific purchasing decisions with additional individual-level data and demographic information, allowing estimated substitution patterns and welfare to reflect heterogeneous tastes for characteristics driven by demographics. Doing so would generate more accurate results, as it frees the model from a heavy dependence on idiosyncratic logit error (Petrin, 2002). Finally, our model and data failed to capture some other potentially important determinants, such as ranking system, although some are highly correlated with the observed characteristics in our model. Having an understanding of additional determinants and the mechanism behind the ranking system can further improve the accuracy of demand estimation. We believe that the limitations of our paper open avenues for more research in this area.

Chapter 4. SOCIAL INFLUENCE OF ONLINE TWO-SIDED MARKET - REVIEW-IN-REVIEW

4.1 INTRODUCTION

User-generated content (UGC) is no longer limited to peer, and high-cost, communication of replying and writing comments. Instead, the last few years have witnessed the popularity of a review-in-review (RIR) function in UGC platforms. This function allows content consumers to provide an evaluation and feedback to a primary review, based on their post-reading quality judgment, by simply clicking on phrases such as "like," "dislike," "helpful," and "unhelpful" to add an RIR. All active users, including the content generator, defined as an individual who creates content such as primary ratings and reviews, and content consumers, are informed of the cumulative number of RIRs, which is recorded, calculated, and displayed at the bottom of the main body of the review. Most mainstream review systems and social network platforms, such as Amazon, Facebook, Twitter, and Yelp, incorporate an RIR function. Allowing users to add their positive or negative evaluation of a product review reveals high-quality (from the crowd's perspective) reviews and enables contributors to provide feedback at a lower cost than when writing a comment or reply, leading to a larger sample of contributors and generating a more robust evaluation. Consequently, high-quality reviews might be selected and highlighted by an accumulation of positive RIRs that deem them "helpful," whereas a low-quality or controversial review might be identified with either very few positive RIRs or a larger body of negative RIRs, such as "unhelpful."

Researchers have devoted adequate attention to the formation of RIRs conditioned on characteristics of either primary review content or rater identity and have identified the factors that lead to more RIRs (Sen and Lerman 2007, Ghose and Ipeirotis 2010, Mudambi and Schuff 2010, Pan and Zhang 2011, Shen et al. 2015). However, they overlook the relationship between generation of primary review content and RIRs from the opposite direction, that is, how content generators' behavior can be affected by RIRs. Prior research has shown that content generators' behavior is driven by social benefits, such as peer attention (Hennig-Thurau 2004, Shen et al. 2015). In addition, RIRs provide active users with an accurate measure of peer recognition and social capital by revealing the volume of "helpful" assessments. Moreover, the social capital reflected by RIRs can be directional because the meaning and the sentiment of RIRs are standardized and formatted in different but fixed ways, e.g. "helpful" remarks signal positive social capital, "unhelpful" remarks signal negative one, and text replies signal a mixture of positive and negative social capital. Given that RIRs signal social capital for all of a platform's active users, examining the behavior of content generators without considering RIR-reflected social capital provides an incomplete understanding of RIRs. Specifically, recent research has shown an association between reputational gain through RIRs and the decision about whether to contribute (Shen et al. 2015).

It is possible that social capital-seeking behavior exerts an influence not only on the decision about whether to contribute to content but also, on a much broader spectrum, on the decision about how to contribute to content. Although the literature identifies social capital seeking as a key incentive to voluntarily contribute (Lerner and Tirole 2002, Hennig-Thurau 2004), very few studies address the impact of the maximization of social capital on how raters generate content. By observing the appearance of RIRs of past reviews by other content generators, one can learn about and anticipate the expected arrival rate of RIR based on the characteristics of the content. In this way, content generators can become well-informed and forward-looking, strategically adjusting their behavior to be consistent with the expectation of RIR arrival to maximize their social capital.

Compared to whether to contribute, the question of how to contribute can be measured through a much larger number of dimensions. We focus on rating behavior in the review-generating process because, although this behavior is one of the most classic and fundamental measures of how to contribute to a review system, it has rarely been investigated in a manner that incorporates social capital incentives. We investigate how the rating level is adjusted as a response to an RIR. Specifically, we expect rational content generators to give a rating based not only on perceived quality, which is the classic determinant of rating valence, but also on the expected social capital gain reflected in RIRs, which is formed by observing and learning from past content. Our research interest is the extent to which content generators (also referred to as raters) choose their ratings as responses to the expected arrival of an RIR. To this end, we quantify the effect of looking forward to social capital gain on the rating-generation process on a leading online review website for films and books.

It is challenging to identify the impact of expected social capital. First, it is an expectation instead of a realization of observations shown as data, and this missing information must be constructed properly so that its effects can be identified. Second, it is not ordinal with respect to rating levels, that is, more stars do not necessarily imply higher social capital. This leads to a typical conditional multinomial discrete choice setting with respect to expected social capital gain. The dependent variable, however, exhibits its ordinal property when it plays a role in response to perceived quality. Goods with higher perceived quality, if measurable, will be mapped with a higher level of ratings. From this perspective, rating shows its rank as more stars' being superior to fewer stars. Past research that ignores expected social capital focuses primarily on the ordinal property of ratings, with perceived quality as a function of goods' characteristics, content generators' characteristics, and social influence, which results in the application of an ordered discrete choice model as the appropriate econometric tool (Ying et al. 2006, Moe and Schweidel 2012, Ho et al. 2014, Lee et al. 2015). Unfortunately, this model is both inconsistent and incapable of explicitly explaining the effect of RIRs in our context.

To overcome these challenges, we apply a two-stage model. In the first stage, we mimic individual learning and the expectation-forming behavior of content generators. We divide the sample data into two partitions according to time, with the earlier dataset used as a learning set for individuals who will rate in the later dataset. We apply a flexible model, as a zero-inflated negative binomial sieve model, to mimic learning behavior. In the second stage, we apply the learning to form an expected RIR in the later setting. With an expected RIR, we propose a partially ordinal discrete choice model that accounts for a mixture of the conditional multinomial model and an ordinal model to account for the mixture of incentives from seeking social capital and as reflecting perceived quality. We further allow the extent of the mixture to depend on individuals' specific characteristics, helping us to infer the heterogeneity of the degree of sensitivity to social capital seeking.

As expected, the results indicate that content generators' rating decisions are based not only on the perceived quality of the good itself but also on social capital gain through the format of RIRs. Specifically, we find that the utility function of raters' social capital increases with more positive RIRs, shown as "helpful," but decreases with more negative RIRs, shown as "unhelpful." RIRs without a clear direction for sentiment, such as a text reply, also increase individuals' social capital gain. Moreover, our model validates several factors of the perceived quality score. We find that ratings still signal the quality of the goods because, in general, high-quality products receive higher ratings. Overall, the effect of social influence is consistent with the literature (Li and Hitt 2008, Moe and Trusov 2011, Lee et al. 2015).

One set of interesting metrics that we include in our data is popularity among content generators (shown as volumes of different types of content) and popularity among content consumers (shown as viewers' vote for "want to watch" and "watched"). Our findings show that popularity among content generators improves quality perception, whereas popularity among content consumers has the opposite effect. Moreover, the results show content generators' reaction to the distribution of past ratings: Products with a higher density of bad ratings generally receive lower ratings. By allowing the extent of the mixture to be a function of individual-level social network characteristics, our results show that raters who seek social capital can be characterized as having either extremely low or extremely high levels of outgoing ties, moderate levels of incoming ties, less involvement in fan groups, and moderate to high levels of ratings experience. Fixing the mixture function to be centered at zero, we find that expected social capital gain accounts for 20.26% of the incentive to rate at a particular level, implying that models that fail to account for expected RIRs might lose a significant level of information. To our knowledge, our model is the first to explicitly account for expected social capital gain and perceived quality jointly in rating decisions.

We organize our paper as follows. Section 4.2 presents the literature and our methodological innovation. Section 4.3 includes our research context and model-free findings. Section 4.4 provides a top-down, detailed demonstration of our model. Section 4.5 presents our estimation strategy, our main findings, model comparison and robustness checks. Section 4.6 includes the use of counterfactual analysis to provide several business insights, and Section 4.7 concludes the paper.

4.2 LITERATURE REVIEW

An emerging literature focuses on the role of RIRs in electronic word-of mouth systems. As the secondary rating that reflects content consumers' reaction to the content itself, RIRs play a significant role as a measure of content quality in a specific dimension for which the RIR is designed. In the context in which an RIR is formed as a "helpful" or "unhelpful" rating, researchers have examined key attributes and mechanisms that constitute UGC perceived as helpful, including review valence and length, rater innovativeness (Pan and Zhang 2011), review extremity, review depth (Mudambi and Schuff 2010), product type (Sen and Lerman 2007, Mudambi and Schuff 2010, Pan and Zhang 2011), subjectivity level, readability, spelling errors, the helpfulness of a rater's past reviews (Ghose and Ipeirotis 2011), and rater reputation (Shen et al. 2015). Based on the theory of confirmation bias, Yin et al. (2016) find that reviews with a rating that involves less confirmation bias are perceived as more helpful. They also summarize and resolve contradictory findings about the impact of rating valence on the appearance of the "helpful" designation. Other studies focus on more context-specific dimensions of the information revealed by RIRs, such as persuasiveness (Zhang et al. 2010) and consumer engagement (Lee et al. 2013, Yang et al. 2016). To our knowledge, all of the prior studies concern the role of an RIR as a measurement of quality but overlook the impact of RIRs on content-generation behavior which is the focus of this study.

An extensive literature investigates how social influence, including but not limited to RIRs, affects content generators' decisions for rating or other dimensions of action space. Using survey data, Hennig-Thurau et al. (2004) identify social benefits as a primary motivation for providing reviews about personal product experiences online. A number of studies have measured the online product rating dynamics and find that earlier ratings have an impact on later ratings (Li and Hitt 2008, Moe and Trusov 2011). Schlosser (2005) argues that a rating decision is influenced by peers:

A rater lowers her rating once she observes peers' negative opinions about the product, whereas she gives a more balanced rating if others demonstrate a high level of dissent. Similarly, Moe and Schweidel (2012) find not only that consumers' willingness to generate content is encouraged if the consumed goods are associated with a higher valance of rating and higher review volume but also that active raters exhibit a differentiation strategy and give more negative ratings.

Lee et al. (2015) differentiate the impact of prior ratings based on the identity of prior raters as friends or strangers. They find that ratings by friends always result in a herding effect, which can be alleviated by the introduction of a social network. Ho et al. (2014) investigate the dynamics of rating behavior by explicitly modeling quality disconfirmation and the perception of system credibility. They identify quality disconfirmation, constructed as a discrepancy between realized quality and expected quality, based on prior ratings, as a main driver for posting. One study that is very similar to ours is that of Shen et al. (2015) which shows that when a reviewer ranking system that quantifies reviewers' online reputations exists, online raters are inclined to rate products with high sales but few existing reviews to maximize attention gain from content consumers, measured as the RIR count, while avoiding competition for that attention. In addition, higher-reputation raters are more likely to engage in strategic imitation. Shen et al. (2015), however, do not directly answer our research question of whether raters strategically adjust their rating levels to maximize their social capital gain from RIRs.

Our model is built on an ordered discrete choice model (Peter, 1980), which is widely applied in research on rating behavior (Ying et al. 2006, Moe and Schweidel 2012, Ho et al. 2014, Lee et al. 2015). Specifically, we use this model to map ordinal ratings with a latent score that reflects not only the rater's post-purchase evaluation of the product but also others' opinions (Schlosser 2005, Moe and Schweidel 2012, Lee et al. 2015). Recall, however, that expected RIRs are conditional on rating valence in a non-ordinal way, suggesting a conditional multinomial relationship with rating. This means that the rating falls into any one of a set of categories that cannot be ordered in a meaningful way with regard to social capital. To reconcile the inconsistency between ordinal and non-ordinal properties of dependent variables, and in the spirit of the finite mixture of the discrete choice model (Gupta and Chintagunta 1994, Greene and Hensher 2003), we propose a partially ordinal discrete choice model by applying a finite mixture of an ordered logit model and a conditional multinomial logit model to jointly account for both ordered and non-ordered incentives. To the best of our knowledge, our research is the first to account for the partially ordinal property of discrete choice by applying a mixture model.

4.3 Research Context and Data

4.3.1 *Setting*

We collect data from one of the leading social network service (SNS) websites that allows registered users to record information and create content related to films and books (the term "products" is used to succinctly refer to both films and books). Registration is free, and unregistered users can read ratings and reviews of both films and books. The platform provides content consumers with a large body of generic and user-generated content and gathers a large group of readers and film lovers that includes 53 million registered users.

Registered users, including content consumers and content generators, participate in many forms of activities with distinct levels of involvement. First, they can react to generated content by RIRs, including adding "helpful" or "unhelpful" and writing a short text reply to an existing rating and review. This will be shown as summary statistics for each type of RIR, attached to each rating and review on the main page. Second, they can either vote for the product that they intend to consume or mark the item as fulfilled, allowing a cumulative number of consuming intentions and people who have consumed to be calculated and shown on the product pages. Third, the webpage provides users with an online social network function by adding links to form outbound social ties and recommending the activities and reviews generated by friends whom users follow. The social tie is directed, meaning that raters who are followed will not receive updates of the activities of individuals who follow them. Notably, the exposure of registered users' information is not limited to content generated by friends. Although content generated by friends will be recommended on the main page in a higher position, information generated by non-friends is still accessible through both searching and platform-wide recommendations.

A content generator enjoys the same level of accessibility and degree of freedom with respect to her online activities as those enjoyed by content consumers. In addition, in the role of content generator, an active user can share her experience about a specific product, with a typical action's involving rating the product with a review. The rating is scaled from 1 to 5 stars in an ordered but discrete manner. The more stars a rating has, the higher is the level of satisfaction expressed by the rater. More specifically, an individual can create three types of ratings and reviews: (1) rating without an RIR; (2) rating with abstract reviews that receive only "likes"; and (3) rating with a review that receives all types of RIRs. Based on our research objective to examine the impact of RIRs, we focus on the third type of rating and review because the variety of RIRs has the potential for a more comprehensive investigation.

Given this dual role, content generators' rating decisions are influenced by experience, perceptions, and learning as a content consumer. As a content consumer, individuals receive social influence, such as how others evaluate a product, by viewing past ratings and reviews. With the introduction of RIRs, individuals can view an unlimited number of products, ratings, and reviews by other content generators for products along with past RIRs related to those ratings and reviews.

Note the role of RIRs as revealing the satisfaction of content consumers who evaluate the rating. Satisfied content consumers consistently reciprocate social capital by adding a positive RIR (e.g., "helpful") or at least not adding a negative RIR (e.g., "unhelpful"). Unsatisfied content consumers will, however, react in the opposite manner. Cumulative user experience and exposure when browsing ratings and reviews by others can explain the relationship between a rating and its corresponding recognition from subsequent RIRs.

Once a user has experienced a product and is willing to share her evaluation of that product with others, suggesting a role transition into being a content generator, learning and social influence will be taken into account for decision making. Users will calculate a quality score to map the ordered rating choices. The score is typically based on the perception of quality based on both their own experience and others' attitudes. However, the incentive underlying the ratings provided by content generators who are sensitive to social recognition might not be limited to reflecting quality. In addition, content generators will provide an appropriate rating valence to satisfy content consumers. Once the relationship between the characteristics of reviews and their corresponding RIRs has been determined, a social capital-seeking individual will choose the rating valence to maximize expected positive RIRs and minimize expected negative RIRs to maximize the expected social recognition utility.

4.4 DATA

We collect a random sample of 5,609 ratings and reviews generated on the webpage during a 180-day window. For each of those ratings and reviews, we also collect RIRs for each day in the subsequent 180 days and three pieces of information, as summarized below.

Ratings and Reviews: We record the rating valence, measured on a discrete 1- to 5-star scale, which includes both the content generators' action that we want to measure and the dependent variable that we will use in the model.

Products: For each product (i.e., books and films), the platform provides both generic product information and user-generated information. We collect information such as product characteristics for each of the ratings and reviews, including product type (book or film), measurement of product quality, the number of posted pictures displayed on the page, the number of individuals who vote for usage intention, the number of individuals who have experienced the product, and the cumulative distribution of the product's ratings, including both volume and valence. For each rating and review, the corresponding product characteristics are captured immediately before the rating and review are posted, reflecting timely information that influences the rating decision.

In terms of product quality, we use an index similar to an IMDb score, calculated as a weighted rank using a Bayesian method with inputs of not only past ratings and volume but also of time and rater identity, thus consistently reflecting product quality. Although this score has partial overlapping information with ratings distribution, it represents not only the best one-dimensional measurement of quality for researchers but also the authoritative and objective aspects of measurements for platform users. In addition, both the nonlinear calculation of the score and a moderate correlation between the score and measures of ratings distribution alleviate collinearity concerns.

With respect to the distribution of past ratings, we individually calculate the percentages of 1- to 5-star ratings to provide a comprehensive measure. In addition, we divide the ratings volume into three groups based on rating type. Recall that, in our context, there are three types of ratings.

This categorization enables us to identify the social influence of popularity from different types of content generators separately. Moreover, the number of individuals who vote for usage intention and the number of individuals who have experienced the product quantify popularity among all active users, who are mostly content consumers. Those variables further and more precisely decompose the social influence from different types of users.

Raters' Social Network Characteristics: With all else constant, content generators with different levels of social status on the platform could behave quite differently. To account for content generators' heterogeneity, we capture each content generator's social network characteristics. We subscript each individual's number of incoming ties and outgoing ties to identify individuals who are the "superstars" of the platform and who seek social recognition. Typically, a user with greater interest in the products, and who is eager to receive more specialized information, is a member of a larger number of fan clubs. Therefore, we collect the number of fan clubs attended by each rater to approximate the depth of expertise level and seniority as a content consumer. Further, we measure the cumulative number of reviews written by each of the raters to quantify her seniority as a content generator.

Review-in-Review (RIR): For each rating and review, a user can express her attitude with three types of RIRs. A positive RIR, shown as adding "helpful," expresses the content consumer's recognition and satisfaction. Content generators feel respected and appreciated when a positive RIR is given, suggesting a positive additional social capital gain obtained by content generators. Conversely, a negative RIR, shown as adding "unhelpful," might depress content generators because their work is not evaluated positively by content consumers, suggesting a negative social capital gain. The third type of RIR is a text reply. Text replies after one review might be neutral, for example, an advertisement for a non-correlated product or neutral-sentiment daily conversation;

however, a review also includes replies with positive social capital, such as appreciation and praise, or negative social capital, such as criticism and condemnation. Therefore, the tone of the text replies can be mixed.

In general, we collect most of the directly visible quantities available across product pages, user pages, and review pages. This information also accounts for the majority of information that the website provides to visitors. In Table 4.10, we exhibit summary statistics for ratings and reviews, products, and raters' social network characteristics. Except for rating valence, average quality score, the dummy variable for movie, and percentage of past ratings, all of the variables are count variables with positive skewness, shown as a long tail on the right, and the mean as larger than the median. In addition, the distributions of the daily RIR count, including counts of helpful or unhelpful and text replies, exhibit zero inflation and dispersion.

Table 4.10. Summary statistics of collected variables

Variable	Mean	Std. dev.	25% quantile	50% quantile	75% quantile	Min, Max
Rating Valence	3.79	0.95	3	4	4	1, 5
Avg Quality	6.84	2.31	6.30	7.50	8.30	0, 9.80
Is Movie Review	0.19	0.40	0	0	0	0, 1
Picture	177.20	268.33	0	66	277	0, 1824
Rating Volume	26886.00	48618.65	537	5602	29975	0, 482356
ShortRev Volume	9604.00	17232.07	142	1816	11310	0, 155622
LongRev Volume	308.00	783.65	11	69	298	0, 11623
Watched	33545.00	63054.02	650	6681	35209	0, 666389
Watching Intention	12432.00	17404.34	860	4818	16401	0, 91398
1star	0.04	0.11	0.00	0.01	0.03	0, 0.94
2star	0.06	0.09	0.01	0.03	0.08	0, 1.00
3star	0.25	0.15	0.13	0.24	0.36	0, 1.00
4star	0.34	1.53	0.24	0.38	0.46	0, 1.00
Outgoing Ties	98.07	226.03	6	21	78	0, 1997
Incoming Ties	947.70	6376.74	10	33	125	0, 80586
Groups	42.74	64.53	3	17	54	0, 395
Total Reviews	165.90	239.29	22	83	184	0, 1348
Helpful	0.03	1.23	0	0	0	0, 539
Unhelpful	0.01	0.15	0	0	0	0, 38
Replies	0.01	0.32	0	0	0	0, 71

Figure 4.1 presents a closer examination of the unconditional distribution of daily RIRs. In the first row is the overall distribution. The histogram chart confirms zero inflation for all three types of RIRs by showing the dominating volume of daily RIRs at the level of zero. To have a direct impression of RIRs at non-zero value, we also provide histograms for three types of RIRs in the interval between 1 and 20 in the second row. The histograms validate the over-dispersion of the distribution of count data.

Figure 4.1. Distribution of volume of RIR per period

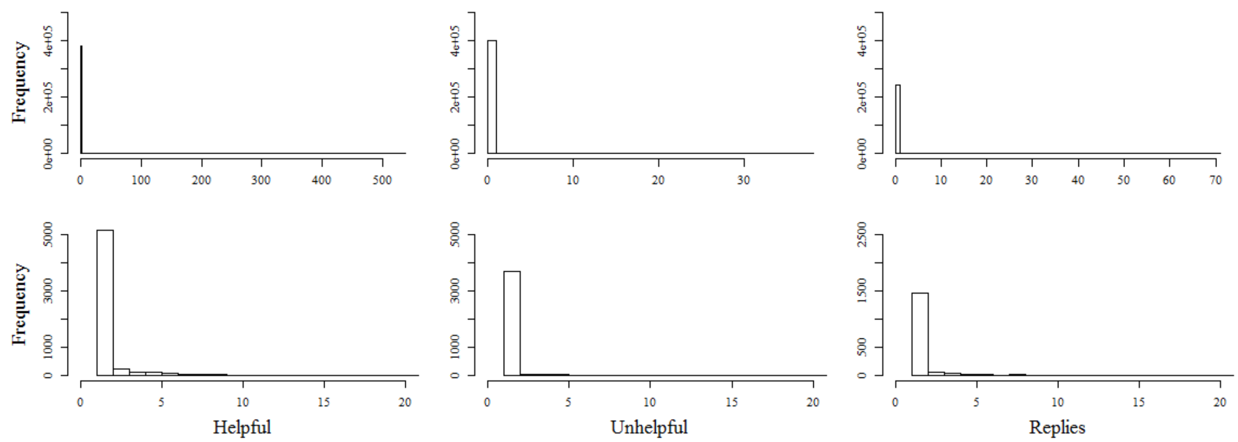

Figure 4.2. Average number of RIR of each rating post per day

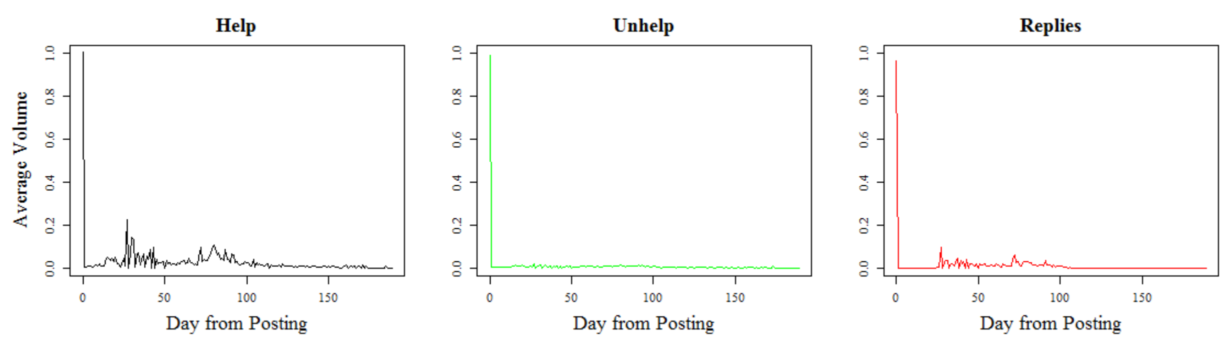

Intuitively, daily RIR volume also is associated with time lapse. Ratings posts will capture the majority of RIRs at the beginning of the posting, leaving the rest to diminish with time. Figure 4.2 presents the longitudinal RIR changes, which confirm our intuition. The average RIR volume is approximately 1 for all three types on the same day when a rating is posted and close to zero afterward, with no apparent linear trend over time. This suggests the importance of modeling expected RIRs conditional on time with a more flexible, non-linear method.

Figure 4.3. Average of 80-day-accumulated RIR of each rating valence and confirmation bias

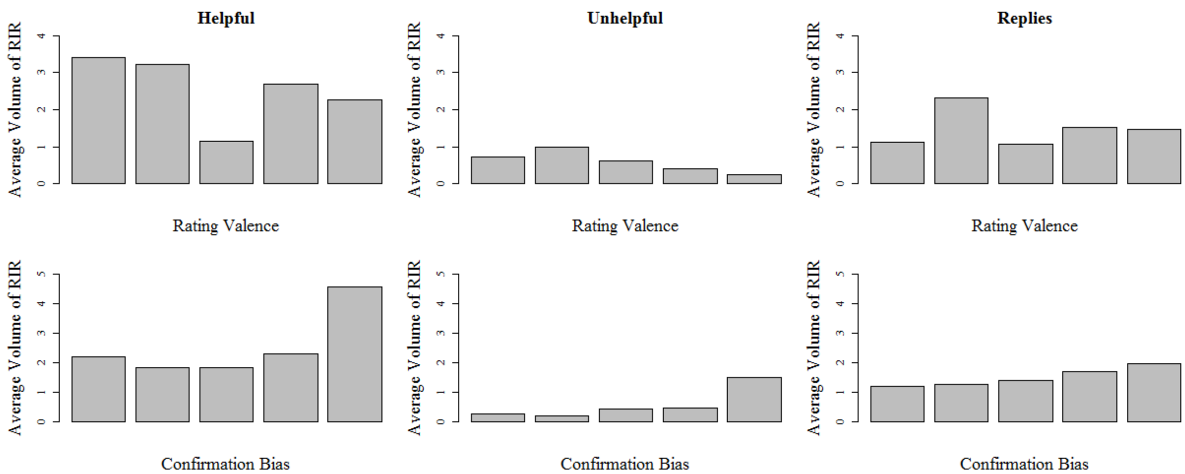

Figure 4.3 displays the association between rating valence and RIRs. The first row shows the average volume of three types of RIRs accumulated in the 80 days after posting, corresponding to the 1- to 5-star rating valence. We find that rating at the median valence level attracts fewer “helpful” RIRs and fewer text replies, which is consistent with our expectation, as median-level ratings are less informative for content consumers. Other than that, it seems that there is no consistent pattern of RIR evolution among the various types of RIRs. Prior research shows that the rating valence alone is insufficient to explain the RIRs, whereas confirmation bias, calculated as the relative rating valence given prior ratings distribution, is a main driver (Yin et al. 2016). Therefore, we include past ratings distribution and compute the distance between rating valence and product quality to represent the confirmation bias. The confirmation bias is rescaled from 0 to 5, with 0 as suggesting the least bias. We divide the confirmation bias level into five groups, from least to most, displaying the average RIR conditional on confirmation bias, as seen in the second row of Figure 4.3. As we expected, the figure shows an association between rating behavior and RIR evolution: A higher level of RIR is accumulated if the rating leads to a higher level of

confirmation bias. This finding is the opposite of that seen in Yin et al. (2010); note, however, that Yin et al. focus on the context of instrumental goods, whereas we focus on books and films. It is entirely possible that, in our context, disconfirmation bias is preferred.

Recall the research objective of assessing the impact of RIRs on rating behavior. To address whether rating behavior is different, given distinct RIRs, we analyze the association from another perspective by calculating the distribution of rating behavior conditional on RIRs. Specifically, we again divide RIR volume into five subgroups by evenly dividing the logarithm of RIR volume, with the leftmost group as representing the lowest volume. Similar to the finding presented in Figure 4.3, the conditional distribution conveys no consistent pattern, as the average ratings are quite close together. Nevertheless, the conditional distribution of confirmation bias exhibits an increasing trend from Level 1 to Level 4 and a decreasing pattern in Level 5, suggesting that, in general, an RIR encourages confirmation bias from the minimal to moderate level while discouraging confirmation bias at a higher level.

Figure 4.4. Average of rating valence and confirmation bias on 80-day-accumulated RIR

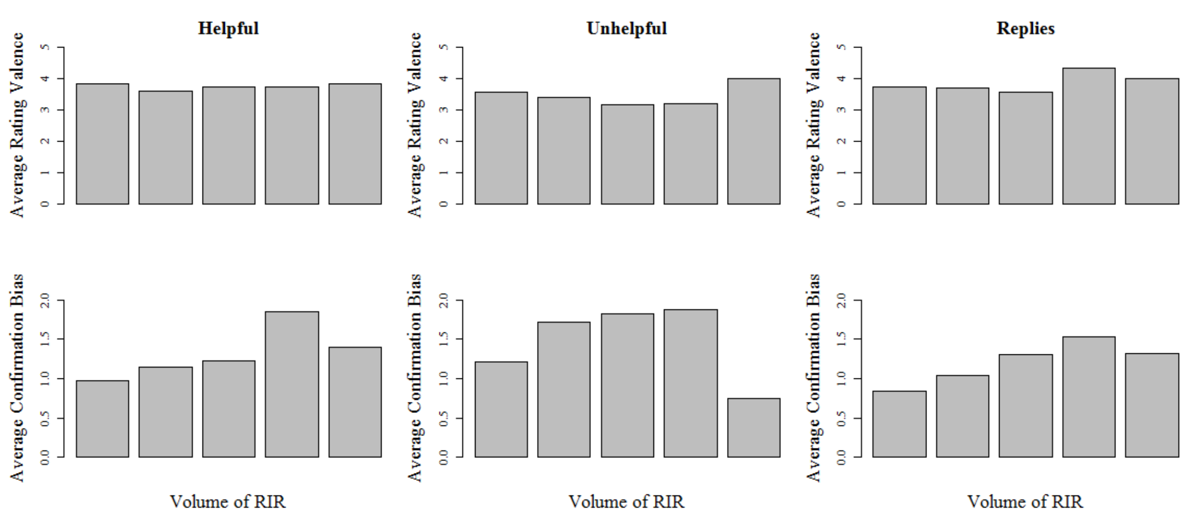

The model-free evidence suggests that the association between rating behavior and RIRs is inadequate to address how RIRs causally affect rating decisions. In Figure 4.4, first, we use the realization of RIRs to visualize rating behavior conditional on realized RIRs instead of expectations. Note that the realization of RIRs is formed after the rating decision. It is counterintuitive to assume that rating behavior is driven by subsequent realization of RIRs due to reversed order. Indeed, the key information that influences rating decisions is the expectation of RIRs, which is formed before ratings are realized. Although expectation is derived from the realization that the two might exhibit similarity, they are not exactly the same. In addition, from the perspective of econometric methodology, realized value cannot be used in this case due to endogeneity issues. Second, although model-free evidence is capable of limited dimensions in each presentation, it fails to consider types of RIRs and predictors other than RIRs jointly. Given that the incentives underlying rating decisions involve multiple dimensions of information, a thorough understanding will be formed only when relevant information is counted. To overcome these two shortcomings, we recover the data-generation process by structurally modeling the expectation of RIRs and apply that expectation together with other information—namely, perceived quality of goods and social network characteristics of content generators—to provide a more comprehensive empirical investigation.

4.5 MODEL

We now set up and estimate a model of rating decisions. First, we explain the model setup of a partially ordinal discrete model and discuss its advantages and properties. Second, we specify the parametric model to construct expected social capital. Third, given the expectation of RIRs, we form the latent utility of social capital gain, conditional on forward-looking expectations, and map the probability of the rating valence with expected social capital gain by applying the conditional multinomial logit model. Fourth, we model how users rate a product as an ordinal response that depends on perceived quality. Finally, we discuss how to mix these two incentives, conditional on the social network characteristics, and formulate the ultimate likelihood function of rating decisions.

4.5.1 *Partially Ordinal Discrete Choice*

We consider a user who decides to rate a product $i = 1, \ldots, I$ at an exclusive and discrete response level of $r \in \{1,2,3,4,5\}$. The rating valence is decided by considering two distinct incentives: perceived product quality and expected social capital gain in the form of RIR. Understanding the inconsistency between rating's ordinal component (by perceived quality) and non-ordinal component (by expected social capital gain) is the key challenge in our research. We address this issue by proposing a partially ordinal discrete choice model that enables a joint and explicit estimation of the impacts of two incentives, with part of the incentive for rating incurred by perceived quality's maintaining an ordinal property, whereas the rest of the incentive incurred by expected social capital gain's failing to do so. Specifically, given the exclusiveness and exhaustiveness of decision r, we use P_{ir} to represent a rater's probability of giving product i a rating valence of r, where $\sum_{r=1}^{5} P_{ir} = 1$ for $i = 1, \ldots, I$. S_{ir} represents the probability that the rater for review i will rate r conditional on the expected social capital gain, and Q_{ir} represents the probability that rater will rate r conditional on the perceived quality of a product. Note that exclusiveness and exhaustiveness hold for S_{ir} and Q_{ir}. To accommodate two incentives that might follow different distributions, we introduce w_i as the individual specific sensitivity weight to perceived quality conditional on individual social network characteristics. Because only the relative weight between two incentives matters, we normalize the sum of weight for perceived quality and weight for expected social capital gain as 1, bounding w_i in the interval of $[0,1]$. Apparently, $1 - w_i$ is the sensitivity weight for expected social capital; higher sensitivity weight for perceived quality indicates lower sensitivity weight for expected social capital gain. Users rate in review i the valence of r by mixing S_{ir} and Q_{ir} as follows:

$$P_{ir} = (1 - w_i)S_{ir} + w_i Q_{ir} \text{ for } r \in \{1,2,3,4,5\}. \tag{4.42}$$

Statistically, P_{ir} follows a mixed distribution of S_{ir} and Q_{ir} with a degree of w_i. Constraining w_i in $[0,1]$ holds the exclusiveness and exhaustiveness property at P_{ir}. Given fixed values of S_{ir} and Q_{ir}, P_{ir} is the average of S_{ir} and Q_{ir} weighted by w_i. Modeling in this manner confirms the intuition that individuals are more likely to rate r if S_{ir} and Q_{ir} are large. In addition, our model shows its strength when considering the dominance of incentives, allowing more flexibility in interpretation. For example, individuals who are sensitive to expected social capital but not to perceived quality would rate r when S_{ir} is large, even if Q_{ir} is small.

4.5.2 *Stage 1: Expectation of RIR Arrival*

A social capital-seeking content generator makes a rating decision by forming an expectation of social capital gain, providing a rating valence to maximize that gain. Given our research focus on RIRs, we model social capital as revealed by three types of RIRs: (1) a positive RIR, indicated by the number of "helpful" RIRs; (2) a negative RIR, by the number of "unhelpful" RIRs; and (3) a mixed RIR, by the number of text replies.

Recall that, although the exact arrival of an RIR is revealed after the rating decision, an individual forms the expected number of RIRs before the decision. To form an expectation, an individual learns the association between RIRs and review characteristics by observing the past appearance of RIRs of ratings and reviews and perceives information about review quality, product characteristics, and rater's status on the website. Based on association rules learned from past reviews, expectations of an RIR for the current review and rating can be formed, conditional on information about its rating valence and review quality, product characteristics, and rater characteristics.

We apply a statistical model to mimic a content generator's learning from past reviews and ratings, along with her expectation-forming process. Specifically, to mimic learning from past reviews and ratings, we divide the data into two parts: The first part includes data with earlier time stamps, defined as past reviews and ratings, whereas the second includes data for later periods, defined as the current review and rating. We use past reviews and ratings to train the learning of RIR arrival rules, mimicking content generators' learning about RIR appearance from their past content. Because this learning is unobserved and unrestricted for the model assumption, to relax restrictions on the rules that govern the appearance of RIRs, we apply the idea of a sieve model to form the conditional expectation of RIRs based on flexible linear forms of the information that is given, including rating valence, review quality, product characteristics and rater characteristics. Due to the limitation of space and that the sieve model is a well-documented predictive method, we refer interested readers to the appendix Section A1, where we discuss the construction of sieve model in our setting in details.

4.5.3 *Stage 2, Part 1: Social Capital Maximization*

We model expected social capital gain u_{ir} when a valence of r is given to rating i. Given that RIR reveals the social capital for a particular rating, we model expected social capital to be a function of expected RIRs. Specifically, we assume that expected social capital is driven by an expected and discounted accumulation of RIRs from the posting time period until T periods later, where T is large enough that the discounting factor will drive its impact close to zero. The log transformation of expected and discounted RIRs is taken to return a better-fitting performance. Moreover, we add a quadratic form of expected and discounted RIRs to allow more freedom fitting, such as diminishing returns. Applying this to all types of RIRs (including "helpful," "unhelpful," and "replies") and using an additive linear form with a constant representing valence-specific cost, the expected social capital gain given a valence of r for rating i follows the function set forth below:

$$u^*_{ir} = \beta_{s1}\ln\left(\sum_{t=0}^{T}\delta^t p_{irt}+a\right)+\beta_{s2}\ln\left(\sum_{t=0}^{T}\delta^t d_{irt}+a\right)+\beta_{s3}\ln\left(\sum_{t=0}^{T}\delta^t n_{irt}+a\right)$$
$$+\beta_{s4}\left(\ln\left(\sum_{t=0}^{T}\delta^t p_{irt}+a\right)\right)^2+\beta_{s5}\left(\ln\left(\sum_{t=0}^{T}\delta^t d_{irt}+a\right)\right)^2+\beta_{s6}\left(\ln\left(\sum_{t=0}^{T}\delta^t n_{irt}+a\right)\right)^2+c_r+\varepsilon_{sr}.$$

(4.43)

p_{irt} is the expected number of positive RIRs for rating i with valence r collected t days after posting; d_{irt} is the expected number of negative RIRs; and n_{irt} is the expected number of mixed RIRs. δ is the discounting factor, which is fixed at 0.95 in our main model, and a is a positive constant number that avoids a log transformation of zero, which is fixed at 1 in our main model. Given both that the expected RIR is close enough to zero when T is 80 and that the discounting factor to the power of 80 is close enough to zero, we fix T at 80. c_r is a constant parameter to recover the average cost of posting a review with a valence of r. β_{s1}, β_{s2} and β_{s3} are coefficients for positive, negative, and neutral forms of expected and discounted RIRs, which are conditional on a valence of r, and $\beta_{s4}, \beta_{s5}, \beta_{s6}$ are coefficients for the quadratic form of three types of expected and discounted RIRs. To identify the discrete choice model, we fix c_5 at 0.

Social capital-seeking raters form expectations of social capital conditional on rating valence r, which varies from 1 to 5 stars. They further compare conditional expected social capital and choose $r = j$ if rating valence j incurs the maximal value of expected social capital u^*_{ij} compared to the expected social capital gain of rating valence other than j, which is denoted as u^*_{ir}.

Assuming random error follows Type 1 extreme value distribution, we derive a closed-form probability that an individual will rate j based on the expectation of social capital gain as:

$$S_{ij} = \Pr\left(\arg\max\left(u^*_{ir}\right) = j\right) = \Pr\left(u^*_{ij} > u^*_{ir}, r \neq j\right) = \frac{\exp\left(u^*_{ij}\right)}{\sum_{r \neq j, r\in\{1,2,3,4,5\}} \exp\left(u^*_{ir}\right)}, \qquad (4.44)$$

where $j \in \{1,2,3,4,5\}$. S_{ij} is the probability of rating j conditional on the expected social capital gain.

4.5.4 *Stage 2, Part 2: Perceived Quality Mapping*

We model the probability that the rater for rating i will rate j conditional on perceived product quality. To quantify perceived product quality, we model it as a latent score whose continuous space can be further divided into several categorical but ordinal rating valence levels. Prior studies show that other than objective quality, perception of quality also would be influenced by factors such as product types (Mudambi and Schuff 2010) and previous review ratings (Moe and Trusov 2011). Therefore, we model perceived quality score q^*_i as a weighted summation of quality indexing of the goods, which, to our knowledge, is the best approximation to objective quality that we can find, and other relevant factors that are identified in past studies. Coefficients play a role as weights: Sign indicates a positive or negative impact, and scale indicates the magnitude of impact. The detailed specification is based on:

$$\begin{aligned} q^*_i = {} & \beta_{q1} pq_i + \beta_{q2} I(pc_i) + \beta_{q3}\ln(pp_i+1) + \beta_{q4}\ln(pr_i+1) + \beta_{q5}\ln(ps_i+1) + \beta_{q6}\ln(pl_i+1) \\ & + \beta_{q7}\ln(pw_i+1) + \beta_{q8}\ln(pt_i+1) + \beta_{q9} p1_i + \beta_{q10} p2_i + \beta_{q11} p3_i + \beta_{q12} p4_i + \varepsilon_{qi}, \end{aligned} \qquad (4.45)$$

where the following product-specific characteristics are included: (a) pq_i quality index of product calculated by the website on a scale of 0 to 10; (b) $I(pc_i)$ indicator of product type, with 1 as representing a film and 0 as representing a book; (c) $\ln(pp_i + 1)$ log of one plus the number of product pictures posted on the product page, capturing popularity among picture posters; (d) $\ln(pr_i + 1)$ log of one plus the total number of pure ratings without RIRs when the rater is posting, capturing popularity among content generators who contribute to Type 1 ratings; (e) $\ln(ps_i + 1)$ is the log of one plus the number of abstract ratings, capturing Type 2 ratings' popularity among content generators; (f) $\ln(pl_i + 1)$ log of one plus the total number rating with reviews by other raters for the same type of goods, capturing popularity among the type of content generators who contribute to Type 3 ratings; (g) $\ln(pw_i + 1)$ log of one plus the reported cumulative number of watching/reading experiences for the product; (h) $\ln(pt_i + 1)$ log of one plus the reported cumulative number of watching/reading intention; (i) $p1_i$ percentage of reviews that rate the product with 1 star; (j) $p2_i$ percentage of reviews that rate the product with 2 stars; (k) $p3_i$ percentage of reviews that rate the product with 3 stars; and (l) $p4_i$ percentage of reviews that rate the product with 4 stars. ε_{qi} represents factor(s) of perceived quality that are observed by raters but not by econometricians: (a) to (b) measure subjective characteristics and webpage setting of the product; (b) to (f) quantify popularity of the goods among various types of content generators; (g) and (h) capture overall popularity among all types of users, including content consumers; and (i) to (l) curve the detailed distribution of rating valence generated by past raters.

To transform the continuous latent score into an ordered discrete response, we define four cutoff points, mapping the latent score to the ordered intervals with the following cutoff points:

$$j = \begin{cases} 1, & \text{if } q^*_i < a_1; \\ 2, & \text{if } a_1 < q^*_i < a_2; \\ 3, & \text{if } a_2 < q^*_i < a_3; \\ 4, & \text{if } a_3 < q^*_i < a_4; \\ 5, & \text{if } a_4 < q^*_i, \end{cases} \tag{4.46}$$

where $a_1 < a_2 < a_3 < a_4$. Assuming the error terms follow a Type 1 extreme distribution, the following is a closed-form solution of the probability of rating valence *j* as being provided conditional on perceived quality;

$$\begin{aligned} Q_{i1} &= \Pr(q^*_i < a_1) = \frac{\exp(a_1 - q_i^*)}{(1+\exp(a_1 - q_i^*))}; \\ Q_{i2} &= \Pr(a_1 < q^*_i < a_2) = \frac{\exp(a_2 - q_i^*)}{(1+\exp(a_2 - q_i^*))} - \frac{\exp(a_1 - q_i^*)}{(1+\exp(a_1 - q_i^*))}; \\ Q_{i3} &= \Pr(a_2 < q^*_i < a_3) = \frac{\exp(a_3 - q_i^*)}{(1+\exp(a_3 - q_i^*))} - \frac{\exp(a_2 - q_i^*)}{(1+\exp(a_2 - q_i^*))}; \\ Q_{i4} &= \Pr(a_3 < q^*_i < a_4) = \frac{\exp(a_4 - q_i^*)}{(1+\exp(a_4 - q_i^*))} - \frac{\exp(a_3 - q_i^*)}{(1+\exp(a_3 - q_i^*))}; \\ Q_{i5} &= \Pr(a_4 < q^*_i) = 1 - \frac{\exp(a_4 - q_i^*)}{(1+\exp(a_4 - q_i^*))}, \end{aligned} \tag{4.47}$$

where Q_{ij} is the probability of rating valence *j* conditional on perceived product quality for rating *i*.

4.5.5 *Stage 2, Part 3: Mixing Two Incentives*

Content generators can exhibit heterogeneous sensitivity to expected social capital gain and perceived quality, suggesting that the relative weight w_i varies conditional on different characteristics of raters, instead of remaining constant. Therefore, to account for heterogeneous sensitivity, we model weight w_i as conditional on raters' social network characteristics in the online community, allowing that certain types of raters exhibit a stronger preference for seeking social capital. Given the constraint that weight w_i lies in the interval of [0,1], we set the weight as the following logit function:

$$
\begin{aligned}
& w_i = w^*_i / (1 + w^*_i), \text{ where} \\
w^*_i = \exp\Big(& \beta_{w1} \ln(uo_i + 1) + \beta_{w2} \ln(uc_i + 1) + \beta_{w3} \ln(ug_i + 1) + \beta_{w4} \ln(ur_i + 1) \\
& + \beta_{w5} \left(\ln(uo_i + 1)\right)^2 + \beta_{w6} \left(\ln(uc_i + 1)\right)^2 + \beta_{w7} \left(\ln(ug_i + 1)\right)^2 + \beta_{w8} \left(\ln(ur_i + 1)\right)^2 + c_w \Big).
\end{aligned} \tag{4.48}
$$

Four variables and their quadratic forms that curve rater's social status in the online community are included: (a) $\ln(uo_i + 1)$ log of one plus the rater's outgoing ties for rating i when an individual generates the content, capturing proactive social activities; (b) $\ln(uc_i + 1)$ log of one plus the rater's incoming ties for rating i when an individual generates the content, capturing social status in the platform; (c) $\ln(ug_i + 1)$ log of one plus the number of interest sub-groups (e.g., a fan group for a particular director, writer, or actor/actress) joined by the rater, capturing the depth of involvement; and (d) $\ln(ur_i + 1)$ log of one plus the cumulative number of Type 3 ratings generated by the rater, capturing seniority as a content generator. Given that all of the variables above are normalized and centered at zero, c_w captures the relative sensitivity of an average content generator. Note that the proportion of likelihood accounted by perceived quality Q_{ij} is monotonically increasing with weight w_i, and w_i is monotonically increasing with w_i^*. A factor having a positive impact on w_i^* results in a higher preference level of Q_{ij}.

Given the probability of a rating valence of j, conditional on perceived quality, a rating valence of j, conditional on expected social capital gain, and the relative weights of those two components, we derive the likelihood function and the log likelihood function as follows:

$$
L = \prod_{i=1}^{I} P_{ij} = \prod_{i=1}^{I} (1 - w_i) S_{ij} + w_i Q_{ij} \text{, and } \quad = \sum_{i=1}^{I} \ln\Big((1 - w_i) S_{ij} + w_i Q_{ij}\Big). \tag{4.49}
$$

4.6 ESTIMATION

4.6.1 *Estimation Specification*

The general estimation scheme follows a two-stage estimation with the maximum likelihood estimation method. In the first-stage estimation, we use data for past reviews, identified as earlier time stamps, to recover parameters for conditional expectations of RIRs. With estimated parameters, we recover expected RIRs in the data for the current rating and review. In Stage 2, given the outcome in Stage 1, we apply maximum likelihood estimation to recover the parameters of perceived quality of goods, expected social capital gains, and the weights of different incentives numerically by using a Newton-type algorithm. To avoid issues that involve the numerical optimizer, we normalize and rescale data by deducting the mean of each variable and dividing by the standard error. Because our research interest does not lie in the formation or structure of the expectation of RIRs, we do not focus on the interpretation of Stage-1 estimation results that incur more than 400 parameters. We report the point estimation, standard error, and Z statistics of Stage 2 of our main model in the following section.

4.6.2 *Estimation Results*

The estimates of the model parameters in the rater's utility function are shown in Table 4.11. In addition to our proposed partially ordinal model with individual heterogeneity (Model (1)), we estimate inferior models as follows: (2) a typical ordinal logit model that considers only perceived quality; (3) a typical conditional multinomial logit model that considers only social capital gain; and (4) a partially ordinal model that incorporates both perceived quality and social capital gain while ignoring raters' heterogeneity. We expect out-performance by our model with regard to goodness of fit and the consistency and robustness of the qualitative findings when comparing those models. Unless otherwise specified, the following discussion focuses on the results from Model (1).

4.6.2.1 Model Parameters

By comparing the effects of positive, negative, and neutral RIRs, we show apparent evidence for a social-capital maximization incentive for the raters. Specifically, we find the main effect for expected positive RIRs to be significantly positive and convex, indicating that raters value positive RIRs; i.e., they prefer to rate at a level that earns the highest number of "thumbs up." The coefficient for quadratic terms is slightly above zero, with a moderate to low level of significance, implying almost linear or slightly increasing marginal utility gain with a (logged) number of positive RIRs. In contrast, a negative sign of coefficient for negative RIRs shows raters' efforts to avoid negative evaluations. The large scale for the coefficient of the quadratic term demonstrates raters' increasing sensitivity to the expected negative RIRs. Interestingly, regardless of positive, negative, or neutral opinions, the number of text replies has a positive impact on raters' social capital gain, with a positive sign of coefficient, coupled with an insignificant and small-in-magnitude coefficient for the quadratic term. In other words, irrespective of the content or sentiment of the text replies, raters are inclined to attract more attention by rating at a level that attracts the most replies. This finding is consistent with that of Shen et al. (2013). Taken together, the expected signs and significance levels, consistent with the intuition from online reviews and mechanisms for gaining social capital, suggest the effective role of RIRs in forming rating decisions. Overlooking this effect might generate a biased estimate when attempting to understand the ratings-generating process. To further measure the quantitative magnitude of expected RIR impact, we calculate the marginal effect of expected RIR impact by simulation-based approach in the appendix Section A3.

Table 4.11. Estimation results and model comparison

Variables	Partially Ordinal with Heterogeneity (1)	Ordinal Logit (2)	Conditional Multinomial (3)	Partially Ordinal without Heterogeneity (4)
Help	4.67***(1.50)		0.23*** (0.08)	10.62** (3.58)
Unhelp	-2.36* (2.03)		-0.49** (0.20)	-10.29* (7.58)
Constant R1	11.39***(2.51)		-2.35*** (0.13)	14.79* (9.60)
Constant R2	15.85***(3.25)		-1.48*** (0.10)	21.71** (9.11)
Constant R3	18.47***(3.65)		0.12* (0.07)	26.00** (10.17)
Constant R4	18.07***(3.40)		0.51*** (0.06)	28.99** (11.84)
Help^2	0.26* (0.21)		-0.14*** (0.02)	-0.95* (0.88)
Unhelp^2	-2.42 (5.12)		-0.41** (0.19)	-10.27* (8.27)
Replies	4.85***(1.37)		1.05*** (0.08)	3.51* (3.11)
Replies^2	0.17 (0.30)		0.07*** (0.02)	0.57 (1.05)
Constant T1	0.74***(0.03)	0.78***(0.03)		0.57*** (0.10)
Constant T2	0.78***(0.03)	0.79***(0.03)		0.94*** (0.06)
Constant T3	0.18* (0.09)	0.27***(0.08)		0.24** (0.11)
Constant T4	-4.15***(0.14)	-4.21***(0.13)		-4.31*** (0.19)
Avg Quality	0.29***(0.09)	0.30***(0.08)		0.29*** (0.07)
Is Movie Review	0.29* (0.19)	0.26* (0.15)		0.45** (0.20)
Picture Volume	0.06* (0.06)	0.05 (0.06)		0.01 (0.06)
Rating Volume	0.10 (0.23)	0.06 (0.20)		0.27* (0.25)
ShortRev Volume	0.08 (0.18)	0.07 (0.16)		0.15 (0.19)
LongRev Volume	0.12* (0.07)	0.20** (0.09)		0.19* (0.11)
Watched	-0.22* (0.19)	-0.35** (0.16)		-0.46** (0.23)
Watching Intention	-0.03 (0.12)	0.03 (0.08)		0.02 (0.08)
1star	-0.54***(0.06)	-0.47***(0.05)		-0.63*** (0.07)
2star	-0.34***(0.05)	-0.37***(0.05)		-0.43*** (0.04)
3star	-0.53***(0.05)	-0.51***(0.05)		-0.57*** (0.05)
4star	-0.13** (0.05)	-0.11** (0.05)		-0.18*** (0.05)
Constant P	0.35 (0.70)			1.53*** (0.17)
Outgoing Ties	0.44 (1.04)			
Incoming Ties	0.48 (0.72)			
Groups	5.88** (2.66)			
Total Reviews	-5.46** (2.05)			
Outgoing Ties^2	-1.23* (0.68)			
Incoming Ties^2	13.89* (7.45)			
Groups^2	5.71** (2.54)			
Total Reviews^2	3.14** (1.31)			
McFadden Rsq	0.41			
-Log Likelihood	3313.10	3424.70	3697.60	3381.80
AIC	6696.20	6881.40	7415.20	6817.60
AICc	6697.07	6881.59	7415.28	6818.12
BIC	6747.59	6904.89	7429.88	6857.25

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

The constants in the component of social capital gain provide us with rating preferences when expected social capital gain and perceived quality are controlled. With a 5-star rating normalized at 0, we find that giving a neutral rating incurs the lowest cost, whereas an extreme rating incurs a higher cost, especially when a rating is extremely high. In other words, individuals generally prefer to rate at neutral and slightly positive levels, followed by a slightly negative level. Raters hesitate to rate a film either extremely high or extremely low. Moreover, extremely high ratings rank last in preference order. This finding is consistent with the observation of ratings distribution, suggesting that more critical ratings incur higher costs in raters' latent utility.

Estimates of coefficients in perceived quality exhibit expected signs. All else equal, a higher rating will be given when a product is associated with higher quality, as represented by a higher score. Movie reviews exhibit systematically higher ratings than do books. One possible explanation for this result is that movies' more entertainment-oriented properties result in less critical ratings. In addition, raters' perceptions are influenced by their peers' evaluations. Popularity among other content generators (information generators)—represented by picture volume, rating volume, short review volume, and long review volume—positively affects quality perception, albeit with a moderate to low level of significance. Interestingly, popularity among content consumers (information recipients) decreases potential raters' latent quality score, a result that is particularly significant for the number of consumers who claim to have watched/read the product. Note that, compared to information generators, information recipients are less involved users of the platform. Given a rater's status as a highly involved user, our results show that raters perceive both popularity among users and highly involved users as positive signals of quality, whereas they perceive less-involved users as negative signals of quality. Moreover, we examine how quality perception is influenced by other raters' ratings. Note that the distributions of rating shown as percentages for 1- to 4-star ratings are standardized before estimation. When we interpret a comparison between the magnitudes of effects, the scale of coefficients must be adjusted with a standard deviation for corresponding variables. With the effect of the percentage of 5-star ratings normalized at zero and adjustments made, we estimate the impact of percentages for lower ratings as 1, 2, 3 and 4 stars (estimated and adjusted as -4.99, -3.91, -3.47, and -0.85, respectively) as negative, with magnitudes of impact as descending when ratings approach 5 stars. This suggests that raters' quality perception conforms to ratings by peers; quality perception scores decrease more drastically when a larger number of extremely low ratings are assigned to a particular product.

Finally, we discuss the parameters associated with rater heterogeneity. Specifically, an insignificant estimate for the first-order coefficient and a moderately significant (but low in magnitude) estimate for the second-order coefficient present a mild inverted U-shape effect of outgoing ties. Note that outgoing ties indicate platform users' proactive social activities. Our results imply that "social seekers" and "social dropouts" are more sensitive to social capital gain as compared to average raters. Although it is intuitive to find a "social seeker" who places greater weight on social capital gain, it is counterintuitive why a "social dropout" is socially sensitive. One possible explanation is the "complement effect" caused by a lack of other forms of interaction with other platform users, resulting in users' devoting more attention to social capital gain through RIR.

With respect to incoming ties, our estimate exhibits a U-shape, that is, raters with many incoming ties and extremely low incoming ties are less sensitive to RIRs. Note that fewer incoming ties indicates higher independence in the social network; e.g., a rater's activity will be viewed only by a very few friends in the social network. In this situation, low sensitivity is understandable, as the users might be self-selected to be "cool" and ignorant of others' opinions. Nevertheless, this finding is consistent with the saturation phenomenon, as cyber stars already attract a great deal of attention (and social capital) from incoming ties; therefore, such users are typically less sensitive to gaining additional attention from RIRs. The number of fan groups exhibits a J shape, implying that raters with more group memberships are less sensitive to social capital gain from RIRs and more sensitive to perceived quality. This marginal effect becomes stronger when a rater participates in more groups because individuals who show a high level of participation in fan groups are fanatic content consumers. Those highly involved content consumers are more willing to show their expertise about a certain product and, hence, less likely to be biased by social capital. Moreover, we find that sensitivity tends to increase with more rating experience, whereas it tends to decrease after a certain threshold, shown as an inverted U-shape. The increasing pattern on the left-hand side of the inverted U-shape might be dominated by raters' learning or awareness, whereas the decreasing pattern on the right-hand side is consistent with the saturation effect, given that raters have already accumulated a large body of social capital through past ratings. These findings show the importance of controlling rater heterogeneity, as different raters might behave very distinctively. The intercept suggests that an average rater has a weight of 79.74% for perceived quality and a weight of 20.26% for expected social capital through RIRs, presenting the significance of considering social capital gains as a complement.

4.6.2.2 Model Comparison and Robustness

We now compare our model with three inferior models with respect to goodness of fit and robustness of the qualitative findings. The results show the consistency of the qualitative findings. Comparing the numerical estimates of Models (1) and (2), we find that these two models' sensitivities to components of perceived quality are very similar. Except for the coefficient for the quadratic form of the expected volume of help, Model (3) also exhibits the same signs of estimated parameters as those of the expected RIR in Model (1), albeit on a much smaller scale. The coefficient for the quadratic form of the expected volume of "helpful" is estimated to be negative, indicating a marginal decreasing pattern of the "helpful" volume. Finally, we estimate the coefficients for Model (3) as having the same sign as the corresponding coefficients for Model (1), except (again) for that of quadratic "help." The scales of the coefficients are inflated when individual characteristics are not taken into account.

To compare the relative quality of different models, given the data set, we calculated several typical model selection indexes, e.g., negative log likelihood, AIC, AICc (given that our sample is small), and BIC. Shown as the smallest compared to the alternative models, those statistics suggest the superiority of our model when overfitting is considered. Specifically, the results show a model specification performance in descending order as Model (1), Model (4), Model (2) and Model (3), thus validating the value of additional information as individual characteristics and expected social capital gain, based on the traditional ordinal logit model. Model (2) exhibits better performance than does Model (3) with respect to fitness, which is consistent with the estimated average weights of social capital and perceived quality when these two components are mixed.

Note that we set several fixed specifications in developing the model. To erase potential concerns about those specifications, we run a variety of specifications to access the robustness of our results. Specifically, we alternate the discount factor, the length of future horizon for calculating the expected social capital gain. In addition, we also increase the size for training the stage 1 parameters. The results are consistent with that of Model (1). See Section A2 in the appendix for details.

4.7 POLICY SIMULATION

In this section, we access counterfactual analyses to help to understand the extent to which the RIR quantitatively affects ratings distribution. Our model is analyzed in the context in which the RIR function is imposed. One intuitive method of quantifying the effect of RIRs is to reengineer the ratings distribution in an alternative context without RIR functions. In this way, the effect of RIRs can be quantified by comparing the alternative outcome with the original outcome. Notably, we apply a conditional multinomial discrete choice model to incorporate the effect of RIRs, in which only differences of RIRs among distinct ratings affect raters' preference order for particular ratings. This provides a simple method of eliminating the impact of certain types of RIRs (or all RIRs) by setting the corresponding RIR to be constant across all five alternative choices, e.g., setting positive RIR at zero for all conditions, including both 1- and 5-star ratings. By doing so, raters obtain no additional benefits with respect to RIRs by making a specific rating choice. In other words, raters will always expect to have a constant level of social capital gain, regardless of which rating the rater gives. This is equivalent to switching off the RIR feature.

Figure 4.5. Ratings Distribution with RIR Fully or Partially off vs. Ratings Distribution with RIR on

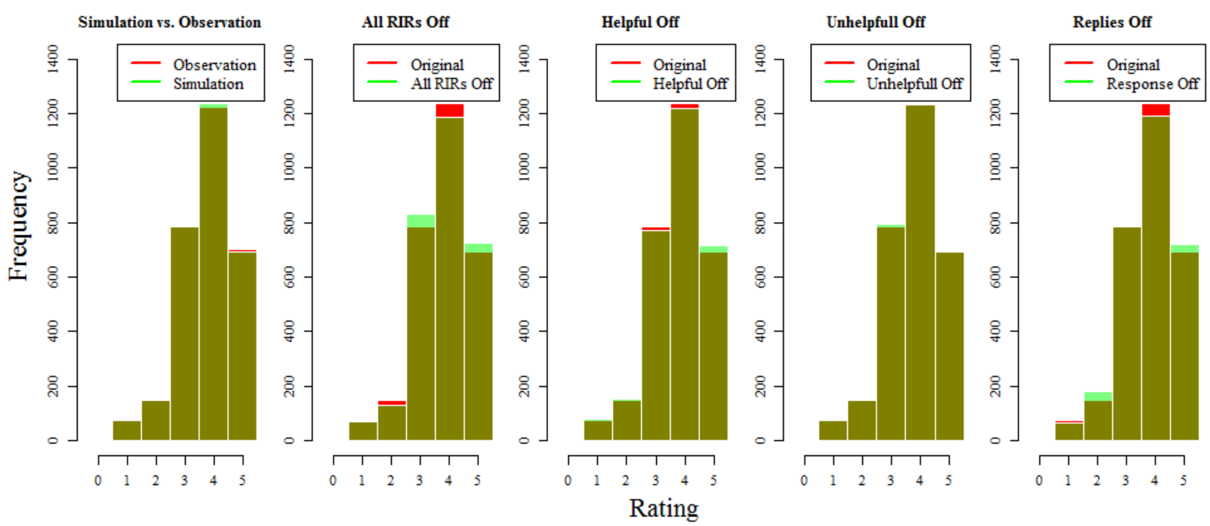

Here, we conduct five specifications of the experiment. The first specification uses our model to simulate ratings distribution with all RIRs turned on. This specification is used to resemble the rating under our observation to check the robustness of our policy simulation. More importantly, recall that we fixed a seed in our simulation. When the seed is used as benchmark to compare with the subsequent simulations, this eliminates noise caused by randomness. In the second specification, we switch off all RIR features, including adding "helpful," "unhelpful," or any text replies to a particular rating. By comparing the rating behavior to the original setting, we can measure how ratings are reshaped with the introduction of RIRs. To further decompose the effect of RIRs into specific sub-types, we further simulate ratings distribution by removing positive RIRs as adding “helpful,” negative RIRs as adding “unhelpful,” and mixed RIRs as adding replies in the final three specifications, respectively. The distributions of ratings under those specifications versus that simulated in the original specification are shown in Figure 4.5.

The results show that our model closely resembles the original distribution that we observed. The simulated distribution exhibits a distribution similar to our observation; e.g., the mean of the observed ratings is 3.7911, and the mean of the simulated ratings is 3.7904. Second, with all RIRs switched off, our model shows that more 3- and 5-star ratings will be given, whereas fewer 2- and 4-star ratings will be given. The average rating is increased to 3.8058, implying that social capital has an overall negative impact on ratings. Third, switching off the option of adding only “helpful” increases 5-star ratings and decreases 2- and 3-star ratings. The change of distribution is quite significant, and the mean of the ratings increases to 3.7959. Fourth, recall that we estimate the parameters of negative RIRs to be smaller in magnitude and moderately significant: The simulation confirms this finding by exhibiting a minimal change of ratings distribution when the option of adding “unhelpful” is switched off, suggesting that this option has a smaller impact than do the other forms of RIRs. The corresponding mean rating is 3.7912. Finally, when the reply option is eliminated from the platform, our simulation suggests a decrease in 1- and 4-star ratings, but an increase in 2- and 5-star ratings. Consequently, the average rating is shown as 3.7901. Comparing the changes associated with a different switching-off policy, it is clear that partially disabling RIRs, e.g., switching off positive RIRs only, has a milder impact on ratings distribution as compared to when RIRs are eliminated, implying that the effect of switching off a particular type of RIR accumulates when the overall system is shut down.

Table 4.12. Transition Matrix When RIR is Switched off Partially or Fully

	1-star	2-star	3-star	4-star	5-star	1-star	2-star	3-star	4-star	5-star
	All RIR off					Helpful off				
1-star	0.19	0.13	0.22	0.38	0.08	0.15	0.07	0.36	0.30	0.12
2-star	0.07	0.17	0.34	0.35	0.07	0.05	0.19	0.38	0.28	0.09
3-star	0.03	0.06	0.43	0.33	0.15	0.04	0.06	0.39	0.37	0.14
4-star	0.01	0.03	0.25	0.50	0.21	0.02	0.04	0.23	0.50	0.21
5-star	0.01	0.02	0.15	0.33	0.49	0.01	0.02	0.13	0.36	0.48
	Unhelpful off					Reply off				
1-star	0.85	0.03	0.04	0.07	0.01	0.28	0.16	0.27	0.22	0.07
2-star	0.01	0.82	0.10	0.05	0.01	0.06	0.32	0.24	0.28	0.10
3-star	0.01	0.02	0.88	0.07	0.03	0.02	0.07	0.49	0.28	0.14
4-star	0.00	0.00	0.05	0.91	0.03	0.01	0.04	0.20	0.55	0.19
5-star	0.00	0.01	0.04	0.04	0.91	0.01	0.02	0.14	0.32	0.51

Although the histograms above present an outcome of ratings redistribution, they might partially overlook the transition process and underestimate the impact of the ratings distribution on the policy change. This is because outgoing transition and incoming transition might offset each other, resulting in the milder change observed in the final distribution; e.g., 50% of 4-star ratings disappear, whereas other ratings change to 4-star and account for 50% of 4-star ratings, leading to no change in the volume of 4-star ratings as the outcome distribution. To avoid underestimation and to accurately quantify the impact of policy change, we further calculate the transition matrix to illustrate the migration of ratings that result from the policy change, shown in Table 4.12. Specifically, we discretely present how the ratings will be redistributed after the policy change for each level from 1 to 5. Each of the table entries is a nonnegative real number that represents a probability, with each row's summing to 1. Each row represents a rating level in the original policy, and each column represents a rating level in the new policy. The percentage of rating j in the new policy conditional on being rated at level i in the original is given as P_{ij} in the matrix, where P_{ij}

is the ith row and jth column element in the matrix P. We calculated the transition matrix for four hypothetical scenarios that we discussed above, in the histogram.

As we expect, the transition shows a much stronger impact of switching off the RIR system. Unlike the marginal distribution change shown in the histogram, the diagonal of the transition matrix suggests that more than half of the ratings change once all RIR features are removed from the platform. Four- and 3-star ratings absorb the majority of transitions. In addition, very low-level ratings are the most significantly affected by the switching-off policy. Further, consistent with our conclusion of parameter estimates, removing the "unhelpful" feature has a weaker effect on the ratings transition than does removing "helpful" or "reply" features, as more than 80% of ratings stay the same even when adding "unhelpful" is disallowed. Conversely, more than half of the ratings are modified when adding "helpful" or a reply is disallowed.

Another angle to look at the quantitative effect of RIRs is to understand how would the ratings be affected if expected RIRs were to behave differently, also known as the marginal effect of expected RIRs. We discuss this scenario jointly with the estimation of the marginal effects of expected RIRs in the appendix Section A3.

4.8 Conclusion and Implications

The RIR system, shown as the secondary rating for generated content, has significantly reshaped the online UGC platform by revealing the social capital of a platform's content generators and viewers. Earlier work has focused on both the formation of RIRs, given generated content and the impact of RIR on subsequent outcomes. However, there has been no research on the impact of RIRs on content generation, which is the major activity of online UGC platforms. Our work, which is based on the assumption that individuals have the ability to learn the formation pattern of RIRs, explicitly addresses this problem.

We frame the decision-making process for rating to incorporate the impact of RIRs by accounting for not only the perceived quality of goods (the classic method of modeling rating decisions) but also the expected social capital gain. Specifically, before making rating decisions, we model raters' learning about expected RIRs by observing past rating content and corresponding RIRs. By mapping the characteristics of content and time with the appearance of RIRs through a flexible sieve model, we assume that raters form expectations of RIRs, conditional on characteristics of content and rating decisions, to be consistent with what they observed in past viewing experience. Then, given consistent expectations of RIRs and the relevant information-signaling quality of the goods, we model rating decisions. Note that decisions driven by perceived quality reveal ordinal information, whereas decisions driven by social capital gain do not. To overcome this challenge, we develop a partially ordinal discrete choice model with heterogeneous weight given to ordinal components, thus allowing raters' decisions to reflect both ordinal and conditional multinomial information.

The results of this model confirm the impact of expected social capital gains: Raters are more likely to give a specific rating that generates more expected positive RIRs, shown as the number of "helpful" tags; fewer expected negative RIRs, shown as the number of "unhelpful" tags; and more expected mixed RIRs, shown as replies that include context. With respect to perceived quality, we find that raters give higher ratings to products with a higher objective quality index, a higher level of popularity among content generators, a lower level of popularity among content consumers, and a lower percentage of low-level ratings. Moreover, the results exhibit a heterogeneous magnitude of the relative impact of expected social capital gain: Raters with extremely low or extremely high levels of outgoing ties, moderate levels of incoming ties, less involvement in fan groups, and a moderate to high ratings experience place relatively more weight on expected social capital gain. An average rater would place more than 20% of the weight on expected social capital gains, implying that expected social capital gain plays an important role in the decision process. To further measure the magnitude of the quantitative impact generated by expected social capital gain, we use simulation methods to experiment with one scenario in which the RIR function is disabled.

Our paper is the first study to incorporate expected social capital gain for content-generation decisions, such as ratings. The newly developed model contributes to the literature by providing a more comprehensive investigation with new dimensions of ratings incentives, further addressing an information problem that potentially has been omitted from studies that have overlooked the social capital perspective. Further analysis also exhibits the superiority of these results with respect to fitting when we compare our results to several other classic models and to robustness when we compare our parameters with those of classic models and alternative specifications.

Our research is subject to several limitations. The first limitation is that our study focuses exclusively on rating decisions, and expected social capital gains might have an impact on many more dimensions of content generators' action space, e.g., the quality and sentiment of content. A model that accounts for all of those dimensions would be more meaningful but very challenging due to the dimensionality problem for dependent variables, given both the difficulty of enumerating all possible dimensions for content generators' action space and the inconsistency among properties of distinct dimensional measurement (e.g., continuous versus discrete and ordinal versus non- or partially ordinal).

Raters can learn about expected social capital gain from all past content until time $t-1$. This implies a distinct training set for each rater in specific time periods. The method we used—setting a threshold time and taking all ratings generated prior to that time—takes the intersection of the individual-rater-specific training set to form the expectation, which could result in less consistency for raters whose t is large, which can be considered the second limitation. Calculating for each rater's expectation based on her corresponding training set would generate more reasonable expectations; however, such an approach involves a very high computational burden for the iterated maximum likelihood estimation for each of those raters.

A third limitation is that we give a very rough measure of the text replies to each rating by simply counting, which pools different types of replies into one group, referred to as a mixture. More detailed findings could be presented through disentangling the mixture by measuring attitudes toward, sentiment about, and the quality of those replies through text mining, as expectations related to distinct types of replies might have different impacts on social capital gain.

Chapter 5. THE MEGA PLAYER OF ONLINE TWO-SIDED MARKET – MEGA APP

5.1 INTRODUCTION

WeChat, seemingly a messaging app, is actually more of a portal, a platform, or even a mobile operating system, depending on one's perspective (Chen, 2015). Launched in 2011, WeChat has one billion registered users and 550 million active users who open the app more than 10 times a day. Usage of the app has contributed $1.76 billion to lifestyle spending and $15.3 billion mobile data consumption in 2014, indicating its mega status in terms of smartphone usage among its users (Cormack, 2015). Industrial anecdotes related to its large scale and user engagement suggest the spillover effects of WeChat. Specifically, its intensive usage might reshape individuals' mobile usage of other apps such that apps with a higher degree of connectivity or functional complementarity to WeChat could achieve high levels of popularity and usage. This effect, however, has not been examined or measured accurately, warranting investigation of the externality of this mega-app.

Recent advancements in app analytics help researchers to understand the usage externality of apps. Ghose and Han (2014) estimate the demand of apps, given their measurable characteristics, and find measurable evidence of the use of in-app purchase design and the removal of in-app advertisements as a means to compete for market share. Other research understands the externality of app demand through special designs for the app marketplace through a rank system, as ranking naturally embeds externality. Carare (2012), who quantitatively measured users' willingness to pay for top-ranked apps, find that it is an additional $4.50 as compared to that of the same unranked app. Garg and Telang (2013) find the "bigger getting bigger" effect, specifically, that the top ranking for paid apps results in 150 times more downloads than the rest of the apps ranked in the top 200 list.

Such research, however, typically focuses on app installation as the measure of usage. Because the post-installation behaviors of users for different apps vary significantly, conditional on the installation of those apps, research is needed to further understand the externality of app usage patterns. Although there is another category of literature in the computer science field that concerns the prediction of post-installation usage patterns (Falaki et al. 2010, Tongaonkar et al. 2013, Xu et al. 2013), such research uncovers only the association rules of app usage patterns and does not provide an interpretation and measure of causality. Thus, such research is insufficient to account for the externality of an app in terms of an economic interpretation.

We address this research gap by estimating the spillover effects of WeChat usage through the use of observational data. This research objective is methodologically challenging for the following reasons. First, given the enormous size of the app market, it is difficult to identify all the apps affected by WeChat. Second, potential endogeneity issues might exist due to the uncertainty of the causal structure. Researchers who fail to account for confounders and the direction of causality might incorrectly take associations as causal effects. Both challenges are extremely difficult to address in the framework of traditional econometrics, when only observational data are available, due to the lack of a causal structure and incomplete information, such as hidden variables.

We propose to integrate a machine learning method with econometrics to identify the spillover effects of WeChat. Specifically, we introduce a Directed Acyclic Graph (DAG) and its unique representation, Completed Partially Directed Acyclic Graph (CPDAG), to characterize the underlying directed causal effect between random variables. Due to the potentially hidden variables that exist behind the observed data, we use a maximal ancestral graph (MAG) and its unique presentation, partial ancestral graph (PAG), to capture causal effects represented by observed variables. We then apply Fast Causal Inference (FCI) and Really Fast Causal Inference (RFCI) algorithms to estimate a PAG uniquely from observational data. Given the estimated PAG, we first identify the adjustment set by two kinds of recently proposed criteria: generalized back-door criterion (GBC) and generalized adjustment criterion (GAC). With the adjustment set and the condition of multivariate normal distribution, we show that the mean causal effects can be estimated quantitatively with a simple econometric linear model.

Our results show that, surprisingly, WeChat has very limited spillover effects on other apps. Only two apps, Taobao and Tencent News, receive positive spillover effects among the Top -50 apps. Our results reveal the true pattern of causality behind the association commonly observed for most of the apps, suggesting that app developers should be reserved about the connection to WeChat, as the spillover effects for most of the other apps might not be as significant as the associations with other apps. In addition, our results emphasize the advantages of using a PAG to estimate causal effects, e.g., uncovering latent confounders (identifying L in $X \leftarrow L \rightarrow Y$ by observing $X \leftrightarrow Y$), avoiding reversed causality (differentiating $X \rightarrow Y$ from $X \leftarrow Y$), and avoiding selection bias (identifying collider in $X \rightarrow Y \leftarrow Z$). We demonstrate these advantages by showing the discrepancy between causal effects encoded in the graph and those estimated with an incorrect interpretation of the causal structure or when the causal structure is unknown.

In our newly introduced method, we use several ways to rigorously evaluate the model performance. First, we test the robustness to additional information by estimating our model, using top-100 frequently used apps and top-300 frequently used apps. Second, we test our model on different weeks, including holiday and non-holiday weeks, and use different samples to ensure its stationarity longitudinally and cross-sectionally in both graphical and quantitative manners. Third, because a PAG needs to perform a conditional independence test, we check the consistency under different specifications of type-1 error levels. The results suggest a high degree of robustness.

To the best of our knowledge, this is the first application paper that integrates the most recent Bayesian network methods as FCI-PAG/RFCI-PAG (PAG estimated by FCI and PAG generated by RFCI correspondingly) and GBC/GAC with econometrics to conduct causal inference. Our research shows the strength of these methods in identifying causal relationships from observational data and suggests the feasibility of determining causal inference when an experimental setting is unavailable or costly. Note that the identification of the causal direction lies in the additional information. This approach also shows its potential in the era of big data, given the ubiquitous availability of additional information. We believe in the potential of the approach to contribute to business analytics area.

We structure our paper as follows. In Section 5.2, we introduce the method; specifically, we explain how to use a graphical model to represent the causal relationship of data. Given the mapping between the data and graph, we then introduce how to recover/learn causal structure from observational data graphically. We then present how to transform the information from the graph into a simple regression that can quantitatively estimate the spillover effects. We include a discussion of the relevant literature and our methods to aid readers' understanding. In Section 5.3, we describe the data that we use in the empirical application, and, in Section 5.4, we present the estimation results. We provide the robustness check in Section 5.5, and, in Section 5.6, we discuss the limitations and provide directions for further research.

5.2 Causal Inference by Graphical Model

A graphical model is an extremely powerful probabilistic tool for modeling the uncertainty within objects, e.g., the conditional dependence structure among random variables. Such a model can provide a clear and effective way to represent a large-scale complex system under mild assumptions. It also can provide a probabilistic inference method within an acceptable time. In addition, the presentation of a graphical model provides an intuitive understanding of the relationship among instances within a system. There are two common types of graphical models: One is Bayesian networks, which are based on directed graph, and the other one is Markov networks, or a Markov random field, which is based on undirected graph. To discover the causal relationships among instances, researchers apply Bayesian networks.

5.2.1 *Graphical Model to Represent Causal Structure*

Bayesian networks were first introduced by Pearl (1982) in the area of artificial intelligence. Later, Pearl developed a probabilistic factorization to represent the causal effect among random variables. Currently, Bayesian networks are a key area of research in machine learning and statistics. For example, as one of the most popular classification methods, Naive Bayes uses ideas of Bayesian networks.

We first introduce the basic definition of a graph. A *graph* can be represented as a pair $G=(V,E)$, where V is a finite non-empty set of vertices, and E is a set of edges formed by linking two different vertices in V, where there is, at most, only one edge between each pair of vertices. In general, there are four types of edges: $\rightarrow$ (directed), $\leftrightarrow$ (bi-directed), $-$ (undirected) and $\rightarrow$ (partially directed). A *partial mixed graph* can contain all four types of edges, while a *directed graph* contains only directed ones, and a *mixed graph* can contain both directed and bi-directed edges. We have a *skeleton* of the graph by ignoring the mark of each edge. If there is an edge between two vertices, then they are *adjacent*. A *path* is a sequence of adjacent vertices. We say that a path is a *directed path* if, for every two adjacent vertices, X_i, X_j, $X_i \rightarrow X_j$ occurs. A *directed cycle* is a directed path from a vertex to itself. A directed graph *G* is called a *DAG* if it does not contain a directed cycle. Given two vertices, *X* and *Y*, if $X \rightarrow Y$, then *X* is a *parent* of *Y*. If there is a path from *X* to *Y*, then *X* is an *ancestor* of *Y*, and *Y* is *descendant* of *X*. Otherwise, *Y* is a *non-descendant* of *X*. A path $< X_i, X_j, X_k >$ is an *unshielded triple* if X_i and X_k are not adjacent. A non-endpoint vertex X_i on a path is a *collider* if the path contains $* \rightarrow X_i \leftarrow *$, where the symbol $*$ represents an arbitrary edge mark. If it is not a collider, then we call it *non-collider* on the path. A *collider path* is a path on which every non-endpoint vertex is a collider.

A *causal Bayesian network* consists of the joint probability distribution of random variables and a directed graph that encodes the causal relationship. Each vertex in *V* represents a random variable. Let *P* be the joint probability distribution of the random variables in *V*, and *G* = (*V, E*) is a DAG; we then define (*G, P*) as a *Bayesian network*. A Bayesian network is a causal Bayesian network if the graph is interpreted causally. The graph and probability are connected through the following two fundamental assumptions (Neapolitan et al., 2004; Pearl, 2011; Scheines, 1997): Markov condition and faithfulness condition.

Markov condition: A DAG and probability *P* satisfies the Markov condition if and only if, for every random variable *X* in *V*, *X* is independent of $V \setminus \{parents(X) \cup Decendant(X)\}$. If the graph satisfies the Markov condition, it means that, for each variable $X \in V$, X is conditionally independent of the set of all its non-descendent *ND*(*X*), given that the set of all its parents *Parents*(X), that is:

$$P(X, ND(X) \mid Parents(X)) = P(X \mid Parents(X)) P(ND(X) \mid Parents(X)) \tag{5.50}$$

This condition not only interprets a DAG as a causal hypothesis but also provides tools for the practice of constructing a Bayesian network by diagnosing such statistical hypothesis testing, which we will discuss later.

Faithfulness condition: If all the conditional independence relations in *P* are entailed by the Markov condition applied to *G*, then it is faithful. When these two assumptions are satisfied, a DAG characterizes conditional independence relationships in *P* via *d-separation* (Spirtes et al., 2000).

A DAG is not fully identifiable. Several DAGs may encode the same conditional independence relation. Those DAGs form a Markov equivalence class that can be uniquely represented by a CPDAG. A CPDAG contains the same skeleton and collider structure as DAG(s). Any edge $X_i \rightarrow X_j$ in a CPDAG means $X_i \rightarrow X_j$ in every DAG in the Markov equivalence class, while an edge $X_i - X_j$ represents uncertainty in the Markov equivalence class, suggesting that both $X_i \rightarrow X_j$ and $X_i \leftarrow X_j$ occur in some DAG(s).

A DAG can represent a causal structure fully in the condition that we have all vertices observed. This condition, however, is barely satisfied when we try to recover the causal structure from data due to the existence of hidden variables or selection variables. Failing to satisfy the condition may cause estimation bias and incorrectly signal a causal relationship. To allow latent variables and selection variables, one can transform the underlying DAG with hidden variables and selection variables into a unique *maximal ancestral graph* (MAG) based only on the observed variables (Richardson and Spirtes, 2002). Recall that a mixed graph has four types of edges. Here, *ancestral graph* is defined as a mixed graph *G* without directed cycles and without almost directed cycles, where *almost directed cycles* occur if $X \leftrightarrow Y$ and $Y \in Ancestor(X)$.

A MAG is characterized by every two non-adjacent vertices X and Y as conditionally independent, given a subset of the remaining observed random variables. In particular, a MAG that contains a tail mark $X -* Y$ means that X is an ancestor of Y *in all DAGs represented by this MAG*. If $X* \to Y$ in M, then, in every DAG represented by M, Y is not an ancestor of X. In addition, the MAG of a causal DAG is called a causal MAG. The conditional independence relationship in a MAG is encoded by m-separation, which is a generalization of d-separation in a DAG (Zhang, 2008). Every pair of two non-adjacent vertices in M are m-separated by a subset of the remaining vertices.

With respect to identification, similar to a DAG, several MAGs may encode the same conditional independence structure and form a Markov equivalent class. Those MAGs could be uniquely represented by a PAG. Like a CPDAG, a PAG has the same skeleton as every MAG in the Markov equivalent class. The relationship between MAGs and a PAG is similar to that between DAGs and a CPDAG. If $X_i -* X_j$ stays constant in every MAG of Markov equivalent class, it will also present as $X_i -* X_j$ in a PAG. If there is an uncertain circle mark in a PAG, such as $X_i \quad -* X_j$, then the Markov equivalent class of MAGs will contain at least one $X_i -* X_j$ and at least one $X_i \leftarrow * X_j$.

5.2.2 *Recovering Causal Structure*

In Section 5.2.1, we showed that the causal structure can be represented by a graphical model. Using a graphical model to conduct causal inference thus consists of two stages. In the first stage, we learn about the causal structure graphically from observational data by recovering a CPDAG (in a hidden-variable-and-selection-variable-free context) or a PAG, which represents all identifiable causal relationships. The second stage involves parameter learning, in which we estimate the causal effects quantitatively based on the graphical structure of Stage 1. We discuss these two steps in detail in the following sections.

5.2.2.1 Stage 1: Recovering Causal Diagram / Learn the Graph

In the literature, there are two approaches to this stage. The first approach is the search-and-score approach that is based on a search procedure and the scoring metric. In this regard, it is to search the best networks by optimizing a predefined scoring metric. Well-known scoring functions include K2-CH metric (Cooper and Herskovits, 1992), chain-based scoring (Kabli et al., 2007), BDeu (Buntine, 1991), Minimum Description Length (Heckerman et al., 1995), and BIC (Schwarz et al., 1978). Because a direct search across all possible graphs is computationally infeasible due to the fact that the number of graphs grows exponentially with the number of random variables, efficient searching or optimizing methods, such as the K2 algorithm (Cooper and Herskovits, 1992), Hill Climbing (Tsamardinos et al., 2006]), Genetic Algorithm (Larrañaga et al., 1996), Simulated Annealing (Wang et al., 2004), Particle Swarm Optimization (Cowie et al., 2007), and Ant Colony Optimization (De Campos and Huete, 2000; Campos et al., 2002), have been proposed to approximate the optimal solutions.

The second approach is the constraint-based learning method that discovers a DAG by testing the conditional independence of random variables. This method is based on conditional dependency among random variables, which is an extension of Pearl's work on Bayesian networks and the Inductive Causation Algorithm proposed in Pearl (1991). For an overview of the constraint-based learning method, please refer to Koller and Friedman (2009) or Scutari and Denis (2014). There are two steps in this method; the first one is the conditional independence test, and the second one is the edge orientation method. In addition, there are some methods, such as the

Max-Min Hill-Climbing (MMHC) algorithm, that combine both of these approaches (Tsamardinos et al., 2006).

Our approach is based on the most fundamental and classic algorithm in the constraint-based learning method; it is a PC algorithm, named for its authors, Peter Spirtes and Clark Glymour, in Spirtes et al. (2000). This algorithm is used to recover a CPDAG when we are free of hidden and selection variables. Starting from a complete graph, in which each node connects with the rest, the PC algorithm gradually removes edges between nodes through a statistical independent test. The algorithm is based on marginally independent tests and then conditional on one vertex's performing conditional independent tests to construct the skeleton and so on. The direction is then added by the algorithm's identifying v-structure and further rules for directions. Kalisch and Bühlmann (2007) have proved the uniform consistency property of the PC algorithm in a high-dimensional setting when the number of variables is a polynomial of the sample size.

The PC algorithm does not work with a MAG or PAG due to hidden and selection variables. To overcome this limitation, an FCI algorithm (Spirtes et al., 2000), which is an improvement of the PC algorithm, is proposed. This algorithm, in addition to the PC algorithm (first-time orientation), incorporates additional steps to remove edges and reorients the graphs based on the PC-oriented collider structure graph. Specifically, the first two steps of the FCI algorithm are almost the same as those of the PC algorithm. In the following two steps, instead of the algorithm's checking all the subsets of the remaining random variables or d-separate set, a superset called *Possible-D-SEP*, as defined Spirtes et al. (2000), can be computed easily. For *G* as a mixed graph, Possible-D-SEP (X_i, X_j) in G is defined as: $X_k \in$ Possible-D-SEP (X_i, X_j) if and only if there is a path *p* between X_i and X_k such that, for every sub-path $< X_m, X_l, X_h >$ of *p*, X_l is a collider on the sub-path in *G*, or $< X_m, X_l, X_h >$ is a triangle of G. It can be shown that the first two steps of the FCI algorithm (or PC algorithm) generate sufficient information to compute a Possible-D-SEP set. Based on the Possible-D-SEP set, the FCI algorithm tests the conditional independence again and reorients the graph based on an updated skeleton and information on the separation set. In the final step, the algorithm uses the orientation rules described in Zhang (2008) to finalize the graph construction. The FCI algorithm has been shown to have the theoretical guarantee that, under some mild assumptions, the sample version of the FCI algorithm is consistent under the high-dimensional sparse setting (Zhang, 2008).

The learning with Possible-D-SEP sets is computationally demanding, rendering infeasibility when the size of the sets is larger than 25 (Colombo et al., 2012). To overcome this issue, some variants of the FCI algorithm, such as the RFCI algorithm and Conservative-FCI (CFCI) algorithm (Colombo et al., 2012), are proposed to help with large dimensional data. The motivation for using the RFCI algorithm is mainly that it tests a smaller number of variables for conditional independent. As a result, the presence of an edge in RFCI-PAG (PAG estimated by RFCI) has a weaker meaning than that of FCI-PAG (PAG estimated by FCI), and RFCI-PAG is theoretically a super-graph of FCI-PAG. RFCI, however, shows great computational advantage, with tolerable errors, when the dimensions of our data are high.

The CFCI algorithm is similar to the Conservative PC algorithm (CPC) proposed by Ramsey et al. (2012). This algorithm is based on two weaker conditions, “Adjacency-Faithfulness” and “Orientation-Faithfulness,” in contrast to Markov and faithfulness conditions. The algorithm can potentially solve some situations when the transitive cause fails. As noted in Ramsey et al. (2012), however, CPC may not be as informative as the PC algorithm, implying that it might be too conservative to discover information. In fact, there is no complete step for orientation on the “unfaithful” mark. In addition, there is no theoretical superiority to assuming the orientation-faithfulness condition and no theoretical property of the further relaxation in CFCI, given that a PAG already assumes a less restrictive condition. Thus, we use FCI to learn a PAG, or RFCI when large dimensions lead to infeasibility or invalidity of FCI-PAG.

5.2.2.2 Stage 2: Estimating Causal Effects / Learn the Parameter

In the second stage, we estimate the scale of causal effects. This step is equivalent to conducting parameter learning of Bayesian networks in the language of artificial intelligence. Given an estimated graphical causal structure, the intuition when estimating causal effects is to control those non-causal effects, e.g., confounders, to adjust the estimated association to be consistent with causal effects. This adjustment is implemented by covariate adjustment.

The classic approach for covariate adjustment in the context of a DAG is the back-door criterion proposed by Pearl (1993). Specifically, a set of variables Z satisfies the back-door criterion relative to an ordered pair of variables (X, Y) in a DAG if:

1. None of vertices in Z is a descendant of X;
2. Z blocks every path between X and Y that has an arrowhead to X.

If Z satisfies the back-door criterion for a DAG G, we could use it to estimate the causal effect between X and Y in a DAG.

It is a sufficient condition to find a set of variables that adjust causal effects consistently. The back-door criterion is applicable, however, only when there is no hidden or selection variables. Because our context has hidden variables, it is infeasible to apply the classic back-door criterion. Therefore, a more generalized criterion is needed to estimate causal effects in a PAG.

We apply two recently developed generalized criteria to estimate causal effects. Worth noticing is that these criteria are available when there is no selection variable, which is satisfied by our first-stage results. The first criterion is a *generalized back-door criterion* (GBC) proposed by Maathuis et al. (2015). It generalizes the back-door criterion to the concept of *visible edge* introduced by Zhang (2008) as: given a MAG M / PAG P, a directed edge $X \rightarrow Y$ in M / P is visible if there is a vertex Z not adjacent to Y, such that there is an edge between Z and X that is into X, or there is a collider path between Z and X that is into X, and every non-endpoint vertex on the path is a parent of Y. Otherwise $X \rightarrow Y$ is said to be *invisible*.

Visible edges refer to situations in which there cannot be such a hidden confounder between X and Y. With the identification of a visible edge, one can extend the definition of a back-door path from X to Y in a PAG / MAG as a path between X and Y that does not have a visible edge out of X. Particularly in a PAG, it means a path that starts with $X \leftarrow *$, $X \quad - *$, or an invisible edge $X \rightarrow$. Zhang (2008) introduces two more definitions to completely define the GBC. One is a *definite non-collider*, which reduces to a non-collider in a DAG or MAG, but, in a PAG, it rules out the possible circle marks. A *definite status path* refers to a path in a partial mixed graph with all non-endpoint vertices as either a collider or a definite non-collider. Following this definition, all paths in a DAG or MAG must be definite status paths.

The definition of the CBC by Maathuis et al. (2015) is as follows: Let X, Y, and Z be pairwise disjoint sets of vertices in G. Then Z satisfies the GBC relative to ordered (X, Y) if the following two conditions hold:

1. Z does not contain possible descendants of X in G;
2. For every vertex $x \in X$, the remaining set of $Z \cup X$ blocks every definite status back-door path from x to any element of Y in G.

The back-door and GBC criteria are equivalent under the DAG framework for a single-intervention setting. Maathuis et al. (2015) propose a sufficient and necessary condition to find such a set that satisfies the GBC criterion. Because the condition requires a lot of graph knowledge, we do not present the condition here. However, we want to highlight that one could easily find the covariates for adjustment conveniently and feasibly compute the causal effects in the data analysis.

The GBC is a sufficient but unnecessary condition for estimating causal effects. Perkovic et al. (2015) further propose a complete GAC that is necessary and sufficient for all of the four types of diagrams that we discuss. The GAC is based on the concept of *amenability*: If a graph G is *adjustment amenable* relative to (X, Y), then every possibly directed proper path from X to Y in G starts with a visible edge out of X. This concept is similar to the definition of the back-door path, but it is defined only on a possibly directed proper path, which relaxes the requirement of a directed path to that of no arrowhead as pointing to the starting vertex. In addition, a path is *proper* from Set X to Set Y if its first node is in X.

The definition of the GAC given by Perkovic et al. (2015) is as follows: Z satisfies generalized adjustment criterion relative to (X, Y) if:

1. G is an adjustment amenable relative to (X, Y);
2. No element in Z is a possible descendant in G of any W, except X, which lies on a proper possible directed path from X to Y;
3. All proper definite status non-directed paths in G from X to Y are blocked by Z.

It is straightforward that both the GBC and GAC are based on intuition in regard to blocking non-causal paths by conditioning on covariate adjustment. Even though the GAC compensates for the shortcomings of the GBC, as it provides only a sufficient condition for an adjustment set, while the GAC provides a necessary and sufficient condition, the GAC does not provide an easily checkable condition, and, thus, there is no algorithm-perspective construction of an adjustment set based on GAC.

Having covariate adjustment set Z via the GBC and/or GAC, one can estimate the causal effects in a PAG. These effects are attained by the definition of the adjustment criterion whereby the motivation of the GBC or GAC is: the set of variables Z of G satisfies the adjustment criterion relative to (X, Y) if, for any probability density f compatible with G, we have:

$$f(y \mid do(x)) = \begin{cases} f(y \mid x) & \text{if } Z=\varnothing \\ \int_z f(y \mid z,x) f(z) dz = E_z\{f(y \mid z,x)\} & \text{otherwise} \end{cases} \tag{5.51}$$

Here, the "do" operator refers to the intervention operator proposed by Pearl (1995) for calculating causal effects in non-parametric models based on the intervention. Equation 5.51 ensures the identifiability of the estimate of the causal effect between variables by transforming intervention probability into conditional probability so that we can estimate the causal effect based on observational study. Once the adjustment set is found, under the Gaussian distribution assumption, the mean of the causal effect is equivalent to:

$$\frac{\partial}{\partial x} E[Y \mid do(X = x)] \quad , \tag{5.52}$$

$$\text{that is } \frac{\partial}{\partial x} E[Y \mid X = x, Z = z], \tag{5.53}$$

Note that we focus only on the linear causal effect. The formula above simply reduces to an econometric model shown as:

$$y_i = \beta_1' X_i + \beta_2' Z_i + c + \varepsilon_i, \tag{5.54}$$

where y_i represents a single vertex that is causally affected, X_i represents a set of vertices that exerts causal effects, and Z_i represents vertices in the adjustment set. β_1' is the parameter vector that capture the causal effects of X_i, which is the one of interest that is to be estimated.

The reduced Model (5) has consistent interpretation in econometrics. The GAC and GBC suggest that controlling Z_i eliminates the non-causal effects of X_i on y_i, which is equivalent to taking Z_i as a control variable to alleviate confounding factors econometrically. Our approach, however, shows its advantage by pinpointing the correct control variables, instead of choosing them simply by assumptions.

5.3 DATA

We use a unique dataset that records app usage behavior of 600 randomly sampled smartphone users in China. For each, we have one observation of the weekly frequency of clicking on all attainable apps on the main user-interfaces of their smartphones. We collect the data for one non-holiday week, starting February 7, 2015, for the purpose of model estimation.

To check the robustness of our finding with respect to stationarity over time, we additionally collect datasets in the same way but for the time windows of the next two weeks (the weeks of February 14, 2015, and February 21, 2015). Note that these two weeks cover the Spring Festival (Chinese New Year), which is an 11-days-long national holiday. This enables us to test whether the causal effects are, in general stationarity, between holiday and non-holiday times. In addition, to check sampling errors, we collect datasets for another sample of 600 individuals that has no overlap with the original sample, in the same way as we execute the original dataset for the same three weeks. In sum, we have two cross-sectional samples and three time periods for each.

The final data (including data for robustness check) included 1,122 different apps, of which 898 appear in the data for estimation. To help readers to have a better understanding of the app market in China, we list the top-50 frequently used apps in China and note the developer and alliance of each in Table 5.13. It is apparent that the app market is not fragmented, suggesting that major developers, such as Baidu, Alibaba, and Tencent (BAT), dominate the app market.

Table 5.13. App Number, App Name, and Corresponding Developer or Affiliation

App No.	App Name	Developer or Alliance	App No.	App Name	Developer or Alliance
1	WeChat	T***	40	91 Lotto	B
2	T Map	T	41	JD.com	T
3	QQ	T	42	B Search	B
4	T Video	T	46	Ali Pay	A
5	QQ Space	T	48	Wo Music	O
6	Weibo	A*	54	MeiTuan	A
7	Other QQ Product	T	55	B Map	B
9	Voice Control	O	59	Moji Weather	O
10	Didi	T	60	QQ Music	T
13	T News	T	68	Iqiyi	O
14	Sogou Typing	S****	72	Tieba	B
15	QQ Browser	T	74	B Wenku	B
17	Youku Video	A	87	Xunfei Plugin	O
18	Kugou Music	O******	88	Baidu Assistant	B
20	Gaode Map	A	91	ZD Clock HD	O
21	B Category	B**	101	Wangyi News	Y*****
22	UC Browser	A	109	WIFI	O
23	360 Guide	O	132	Sohu News	S
24	TouTiao	O	146	App Store	O
27	Android MKT	O	149	Sohu Video	S
29	MiLiao	O	152	Fun TV	O
32	91Phone Assistant	B	188	App Market	O
33	B Map Plugin	B	196	Coolpad Weather	O
35	Sina News	A	239	Kowo Music	B
39	Taobao	A	332	Momo	A
*A = Alibaba **B = Baidu ***T = Tencent			****S = Sohu *****Y = Wangyi ******O = other or independent developers		

In the data for estimation, the weekly average clicking rates for different apps exhibit a typical long tail, with WeChat's on the very left-hand side, as shown in Figure 5.1 (a). In Figure 5.1 (b), a closer examination of the top-50 frequently used apps listed in Table 5.13 shows that the usage of WeChat (the very left-hand side) is at least two times that of the second most frequently used app, confirming its mega status in app usage. We further check the stationarity by including the dataset for a robustness check and depict the average weekly clicking rates across 1,200 individuals over three weeks in Figure 5.1 (c) and Figure 5.1 (d). A comparison with Figure 5.1 (a) and Figure 5.1 (b) shows similar shapes but fatter tails for their distributions.

Figure 5.2 (a) presents the distribution of WeChat usage, for which the clicking rates are quite skewed, with the majority of clicking rates as less than 5,000, with the maximum above 20,000. The skewness suggests a potential problem if we want to make use of a multivariate normal distribution for the estimation of causal effects. Therefore, we take a log transformation of our data to approximate a multivariate normal distribution. For WeChat, the transformed data are shown in Figure 5.2 (b).

Figure 5.1. App Weekly Average Clicking Rates of Estimation Sample and Pooled Sample

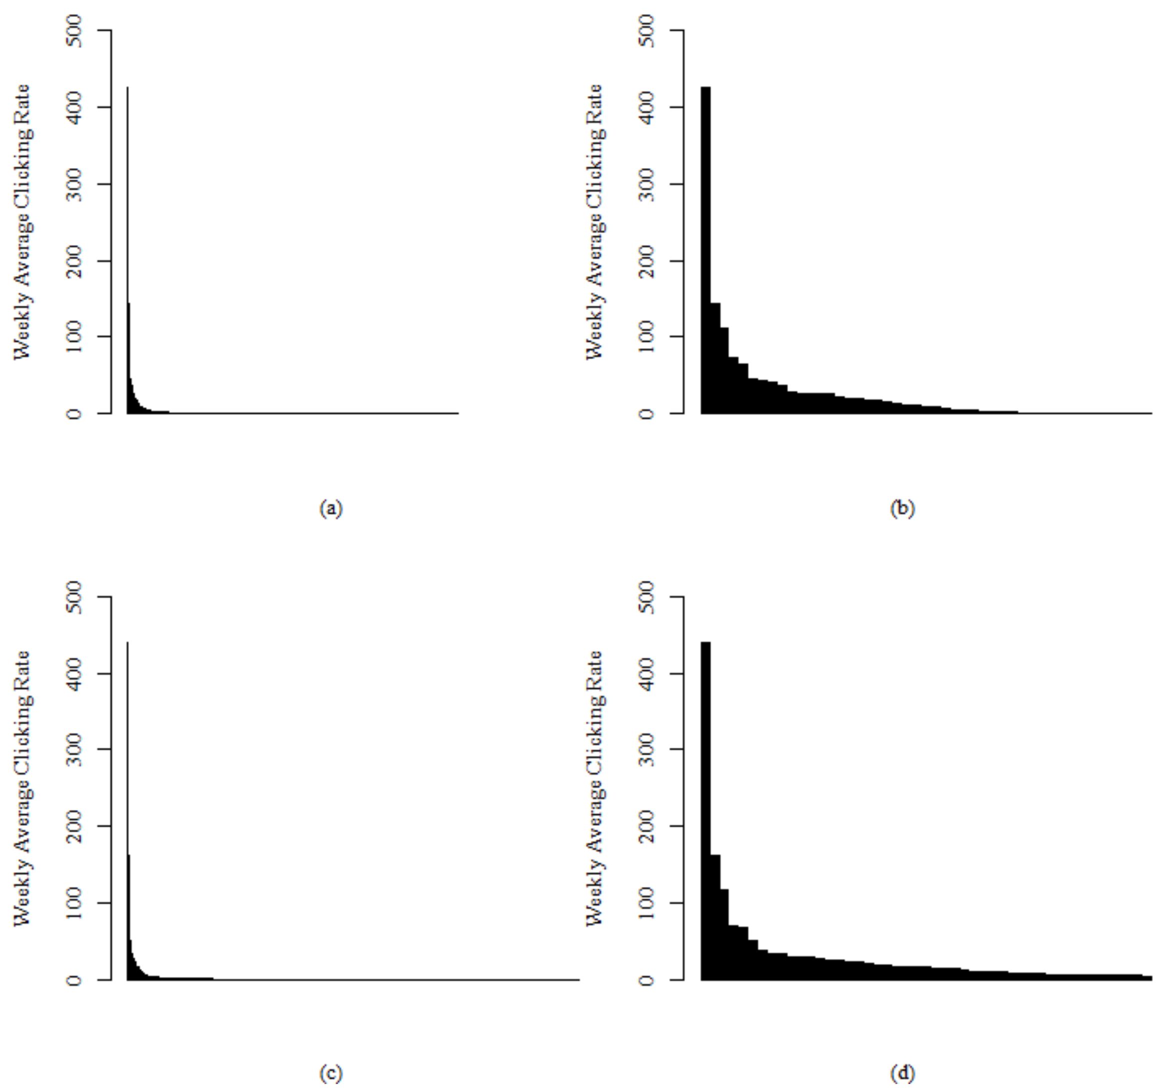

Figure 5.2. Distribution of WeChat Usage

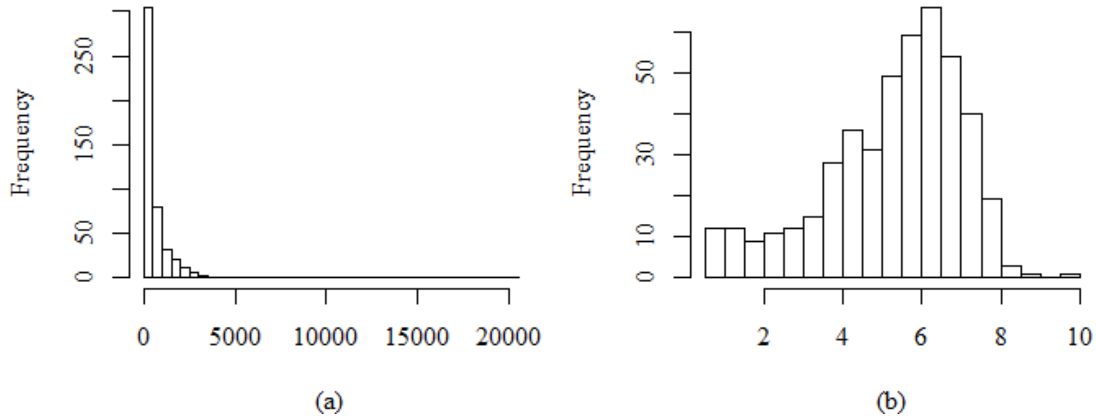

5.4 Estimation Results

Our goal is to capture the causal relationships between different apps, and if there is such a relationship, we hope to estimate the causal effects based on the observational data. We assume that there is (possibly) no directed cyclic graph between apps, which is practical in reality and satisfies the faithfulness condition. Considering that there might be hidden apps behind the data and that selection bias may exist, instead of constructing a CPDAG, we use a PAG to model our data to reduce bias and attain lower variance than would be seen in a CPDAG. In addition, the space of PAGs is smaller than that of CPDAGs, which makes the search more feasible. When the sample is large, the same data with a single PAG can solve a lot of meaningful questions behind the app data, while a CPDAG might give us a different graph structure. Given an estimated PAG, in the second stage, we further quantitatively estimate the causal effects by applying the GAC and GBC to find the valid adjustment set.

We estimate the causal relationship of the top-50 most-used apps only in the main model for following reasons. First, the usage of many rarely used apps exhibits no dependency on the rest. Having a smaller set generates a more concise presentation. Second, those rare apps typically focus on niche markets, which have a less significant impact on the app market as compared to that of top ranked apps. Third, methodologically, (log transformation of) usage of rarely used apps can barely satisfy normal distribution assumptions, which could not only lead to problematic results but also contaminate the results of those frequently used apps. To alleviate concerns about this approach, we extend the set to include more apps for the analysis in Section 5.5.5. Compared to a PAG estimated with an extended set of vertices that includes more apps, PAG estimated with top 50 apps shows that the spillover effects of our focal app, WeChat, are well captured and depicted locally.

We present the results as follows. First, we provide the causal structure of app usage graphically as the PAG that we determined through the FCI algorithm. Second, we measure the spillover effects quantitatively based on the estimated PAG using the GAC and GBC criteria, with econometric interpretation. The quantitative measurement provides further information on the causal effect as positive or negative as well as its strength. Third, to show the value of the graphical model for estimating causal effects, we extend our discussion to cases that are assumed to be estimated without knowing the causal structure from a PAG or with an incorrect adjustment. In those examples, the existence of spillover effects is ruled out by graphical results and interpretation; however, these effects are estimated to be significantly not zero due to the bias of incorrect adjustment.

5.4.1 *Stage 1: Graphical Results*

We present our estimated causal diagram in Figure 5.3. In this diagram, each node shown as a number represents an index of one specific type of app, which is the App Number in Table 5.13. The diagram explicitly displays local causal effects of WeChat (App 1). Note that the edges out of WeChat are visible ($1 \rightarrow 13$ and $1 \rightarrow 39$). This indicates that there are no unobserved confounders behind a direct edge and that each directed edge out of WeChat represents corresponding causal effects explicitly. Specifically, the diagram shows that WeChat has direct spillover effects on two apps: Tencent News (App 13), a news app developed by the same parent company, and Taobao (App 39), the leading shopping platform in China, developed by the Alibaba group. Other than these two apps, WeChat exhibits direct correlations with other QQ products (App 7) and Appstore (App 146), driven by unobserved confounders (as they are connected bi-directly). Figure 5.3 suggests that the correlation between all other apps and WeChat is confounded by hidden variable(s) that are not observed and/or conditionally driven by colliders (observed selection variables) in the data. In sum, the diagram suggests that, even though WeChat dominates smartphone user app use, its direct externality toward other apps is not as strong as we had expected. In fact, it is so limited that only two other apps are affected directly.

The finding suggests that, although associations between WeChat and other focal apps might be found, they are not necessarily explained causally. In fact, for the majority, it is confounders rather than spillover effects from WeChat that explain the association. App developers should be cautious about being deceived by associations when analyzing attribution and collaboration, as the identities of factors that determine the usage of apps might not be the same ones that show the association of usage with the focal app. Given that a connection to such mega apps might incur high costs, our approach provides a tool that allows app developers to visually and directly examine the spillover effects from WeChat and other apps. Our approach provides an understanding that is deeper than that provided by superficial association and helps app developers with decision making with regard to developing collaborations and connections for economic interests.

Worth noticing is that the estimated PAG contributes not only to qualitative but also to quantitative findings. Any node without a (possible) causal path from WeChat is indicated as having no causal effects from WeChat. Therefore, it can be concluded quantitatively that all nodes in Figure 5.3, other than Tencent News and Taobao, receive zero causal effects from WeChat.

Figure 5.3. PAG of (Top 50) App Usage Causal Structure

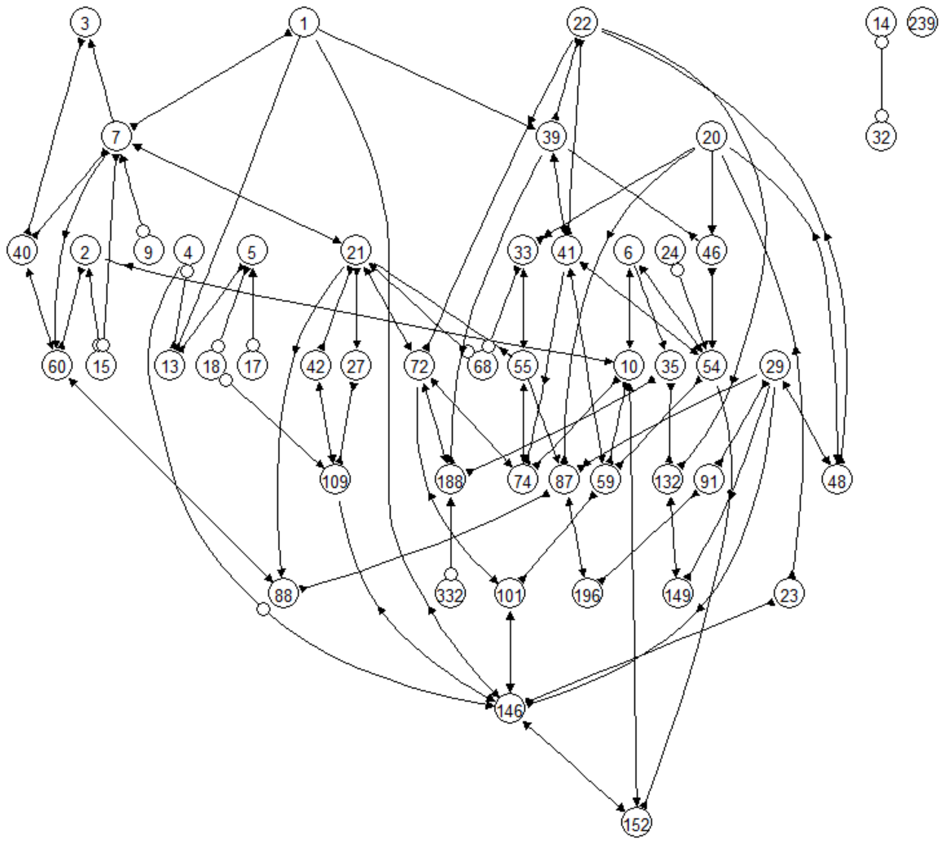

5.4.2 *Stage 2: Quantitative Results*

Given the results for apps that receive zero causal effects from WeChat, however, for apps that receive non-zero spillover effects, we need to estimate the scale of them quantitatively in additional steps. Specifically, to avoid potential biasness due to observed confounders, unobserved confounders, and selection variables, we use the causal structure estimated by the FCI algorithm in Figure 5.3 to adjust non-causal factors, following the GAC and GBC. Figure 5.3 shows that non-causal paths are all blocked by colliders for both Tencent News and Taobao, implying that the adjustment set Z is an empty set, following the GAC or GBC. The model simply reduces to a linear regression with the usage of the focal app, WeChat, as the only independent variable.

Table 5.14 shows that the spillover effects of WeChat are positive for both Tencent News and Taobao. Specifically, for an average user of WeChat, a 10% increment of usage of WeChat leads to 7.25% additional usage of Tencent News and 8.33% more usage of Taobao. This suggests that, as different types of apps are created by the same developer, the functionality of WeChat complements that of Tencent News effectively. WeChat users who are interested in reading news are successfully directed to the news app developed by the same company, indicating one more step to the goal of full service of Tencent. However, the spillover effect on Taobao suggests positive externality to Alibaba, the major competitor of Tencent, given that Tencent has its own online shopping platform and other ecommerce platforms as a strategic alliance. The existence of spillover effects suggests a loss of users with the intention of online shopping, as provided by the competitor.

Table 5.14. Estimation Results

Parameter	Tencent News (13)	Taobao (39)
β_1	0.35***(0.02)	0.40***(0.03)
c	0.20*(0.10)	0.35*(0.14)
Marginal Effects (10% in X)	7.25%	8.33%
Signif. codes:	0 '***' 0.001 '**' 0.01 '*' 0.05	

5.4.3 *Estimates based on Incorrect Adjustments*

The value of a graphical model is not limited to aiding the estimation of causal effects, as shown in Section 5.4.2. Moreover, the estimated causal structure itself encodes enormous interpretable information on causal effects that helps researchers to have an understanding of correctly adjusted causal effects, which would otherwise be incorrectly estimated. In this section, we present several common representative cases in econometric causal inference that appear in our context, including unadjustable latent confounding bias, adjustable latent confounding bias, and endogenous selection. Note that the value of a PAG is not limited to the three cases that we mentioned above. In addition, it can solve over-controlled bias, observed confounding bias, and so on (Elwert, 2013). We skip those issues, however, because those cases do not appear in our context. Further, an incorrect adjustment can happen in any vertices in our data. Due to space limitations, we illustrate only three cases that occur in our data through three representative vertices.

5.4.3.1 Unadjustable Latent Confounding Bias

Based on the interpretation rule of a PAG, a bi-directed edge $A \leftrightarrow B$ suggests that *A* has no causal effects on *B* (due to the arrowhead at *A*), and *B* has no causal effect on *A* (due to the arrowhead at *B*). There is no ancestral relationship between *A* and *B*, but they are adjacent. Therefore, the association between *A* and *B* can be explained only by latent confounder(s) (Kalisch et al. 2012). Because the confounder(s) are unobserved, the confounding bias cannot be adjusted. Therefore, a linear regression model cannot correctly estimate the causal effect between *A* and *B*. A naïve regression of *A* on *B* would induce the confounding bias due to the unobserved confounder.

Table 5.15. Example of Unadjustable Latent Confounding Bias

	Parameter	Other QQ product (7)	App store (146)
Association	β_1	0.45***(0.03)	0.17***(0.02)
	c	1.25***(0.13)	0.02(0.01)
Causal Effects by PAG		0	0
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05			

In our example, unadjustable latent confounding bias exists between the usage of WeChat and that of other QQ products as well as between the usage of WeChat and that of Appstore. The interpretation of a PAG suggests no causal relationship between WeChat and QQ products or Appstore. However, researchers would estimate the causal effect as positively significant if they have no information about the causal structure and mistakenly regard the association as causal effects. We estimate the association and compare it with the causal effect based on a PAG in Table 5.15.

This result shows the methodological advantage of a PAG for estimating causal effects from observational data with hidden confounder(s). Other methods for causal inference alleviate the confounding bias by controlling potential confounding factors, such as propensity score matching. However, such an approach is limited to conditioning on observed confounder(s) only, leading to biased estimation when unobserved confounders exist. The PAG approach, in contrast, infers the existence of an unobserved confounder, which further helps researchers to adjust causal effects correctly.

5.4.3.2 Adjustable Latent Confounding Bias

Latent confounding variables are adjustable when observed intermediate non-collider vertices exist on the causal path from the latent confounder to focal variables. The simplest example is $A \leftrightarrow B \rightarrow C$. In this example, *A* has no causal effect on *B* and *C*. However, *A* and *C* show an association due to a common confounder between *A* and *B*. This confounder exhibits a causal effect on *C* indirectly through *B*. Given that the edge between *B* and *C* is visible because A points to B and B is a non-collider, conditioning on *B* would control the causal effects from the latent confounder to *C*. Therefore, a linear regression of *B* and *C* would adjust the latent confounding bias. If $A \leftrightarrow B \rightarrow C$ is the only unblocked path between *A* and *C*, the regression that suggests zero as the coefficients for *A* can be used as the validation for the bias of the adjustable latent confounding variables.

In our example, one apparent path with latent confounding bias is from WeChat to QQ (App 3), another instant messaging app developed earlier by Tencent, through other QQ products, shown as $1 \leftrightarrow 7 \rightarrow 3$. Note that there is no other unblocked path between WeChat and QQ. The graph suggests that adding usage of other QQ products in an adjustment set Z would control the causal effect from WeChat to QQ. The results in Table 5.16 confirm our expectation by showing the causal effect of App 1 on App 3 to be insignificantly different from 0. The estimation of causal effects without controlling the usage of other QQ products would result in a biased estimation due to failing to adjust for the effect of unobserved confounder(s) between App 1 and App 7.

Table 5.16. Example of Adjustable Latent Confounding Bias

	Parameter	Adjusted	Unadjusted
Association	β_1	0.01(0.02)	0.45***(0.03)
	β_2	0.96***(0.02)	
	c	0.29***(0.07)	1.49***(0.01)
Causal Effects by PAG		(7) has effect on (3)	0
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05			

This finding also shows considerable consistency with recent observations and anecdotes about the relationship between QQ and WeChat, two instant messaging apps by the same developer, from an industry perspective. An industry observer reported that WeChat was designed strategically to differentiate itself from QQ, such that very limited substitution exists (Geekpark, 2013). This observation was confirmed by the CEO of Tencent (ithome, 2013). Individual users would be driven to use these two apps based on different functional needs, such that no direct dependency between these two apps should exist. Other confounder(s), however, might encourage usage of both apps, which would result in association, consistent with our estimation results.

5.4.3.3 Control Over Selection Variable

In the two cases above, we show the potential bias due to failing to control non-causal factors. In econometrics, such cases are typically due to failing to have confounders as valid control variables. This leads to a concern about whether this means that we should have as many control variables as possible to alleviate biasness to the maximal level. In this section, we present a problematic estimation if the control variable is a collider (selection variable), rather than a confounder, on the path. Note that the PAG identifies the role of each node on a path as a collider (or not). This again shows a methodological advantage as compared with models that have an uncertain status of the confounder or collider of each control variable before estimation.

The problem of endogenous selection bias occurs when a collider is added into the adjustment set *Z*. Specifically, conditioning on the common outcome of two variables induces a spurious association between them for at least one value of the collider (Elwert 2013). For example, when we have a PAG shown as $A \leftrightarrow B \leftrightarrow C$, this suggests one possible structure with two latent variables, revealed as $A \leftarrow L_1 \rightarrow B \leftarrow L_2 \rightarrow C$, and that *A* does not have any causal effect on *C* if there is no other path or if all other paths are blocked. However, if we condition on observed vertex *B*, the causal structure will be replaced as $A \leftarrow L_1 - L_2 \rightarrow C$, where *A* is associated with *B* due to the spurious path between L_1 and L_2. Because L_1 and L_2 are unobservable, and thus cannot be added into the adjustment set to block this spurious path, a spurious causal effect will be estimated to represent the endogenous selection bias.

There are many potential examples of endogenous bias if we do not design the adjustment set in the correct way. We take a causal relationship between WeChat and 91 Lotto (App 60), the leading online lotto marketplace in China, as an example. According to the estimated PAG, the causal effect from WeChat to 91 Lotto is 0 because there is no causal path from WeChat to 91 Lotto. However, if we erroneously add usage of other QQ products (App 7) into the adjustment set *Z*, the causal effect from WeChat to 91 Lotto is estimated to be significantly negative, as shown in Table 5.17. This is because conditioning on other QQ products opens a spurious confounding path between WeChat and 91 Lotto, whose confounder is unadjustable ($1 \leftarrow L_1 - L_2 \rightarrow 60$). This example provides important information for researchers: Adding an incorrect control variable risks deteriorating the estimation of causal inference.

Table 5.17. Example of Endogenous Selection Bias

	Parameter	Adjusted
Association	β_1	-0.07**(0.02)
	β_2	0.62***(0.03)
	c	-0.23*(0.10)
Causal Effects by PAG		0
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05		

5.5 Robustness Checks

As presented in this section, we conduct a robustness check to ensure the consistency of the findings and to eliminate potential explanations, such as sampling errors and time-specific factors. Given the nature of the two-stage estimation for causal effect estimation, we first check the consistency of graphical outputs, as discussed in Section 5.5.1, and then check that of quantitative results, as presented in Section 5.5.2.

5.5.1 *Check Graphical Results*

To eliminate concern about sampling errors, we use the alternative sample, in which there are no overlapping individual users. To eliminate the concern about time-specific factors, we collect further data for the next two weeks. Note that two weeks after the time of the original observation is a national holiday. It would imply a high degree of consistency if the spillover effects of WeChat in the original PAG are the same or close to that in the PAG of the holiday. Given the two sets of samples and the three time periods for each, we could estimate six PAGs. For succinct presentation, we draw graphs of only the causal paths from WeChat. Further, we apply RFCI when FCI is infeasible or invalid. The PAGs are displayed in Figure 5.4**.**

As can be seen in Figure 5.4, the causal paths from WeChat are quite consistent for all six samples. All PAGs show direct causal effects on Taobao (App 13) and Tencent News (App 39), suggesting that the causal effects identified in the original sample are robust to different samples and, thus, robust to sampling errors and time-specific factors. The mild discrepancy lies in the PAG of Week 2 and Sample 1, which exhibits an indirect causal effect on App 39 through App 41; the PAG of Week 1 and Sample 2 exhibits an indirect causal effect on App 35 through App 39; and the PAG of Week 2 and Sample 2 exhibits an indirect causal effect on App 39 through App 46. These effects are quite unstable, however, and could be attributed to sampling errors or time-specific factors.

Figure 5.4. PAGs for Two Sets of Individuals and Three Time Periods with Alpha = 0.01

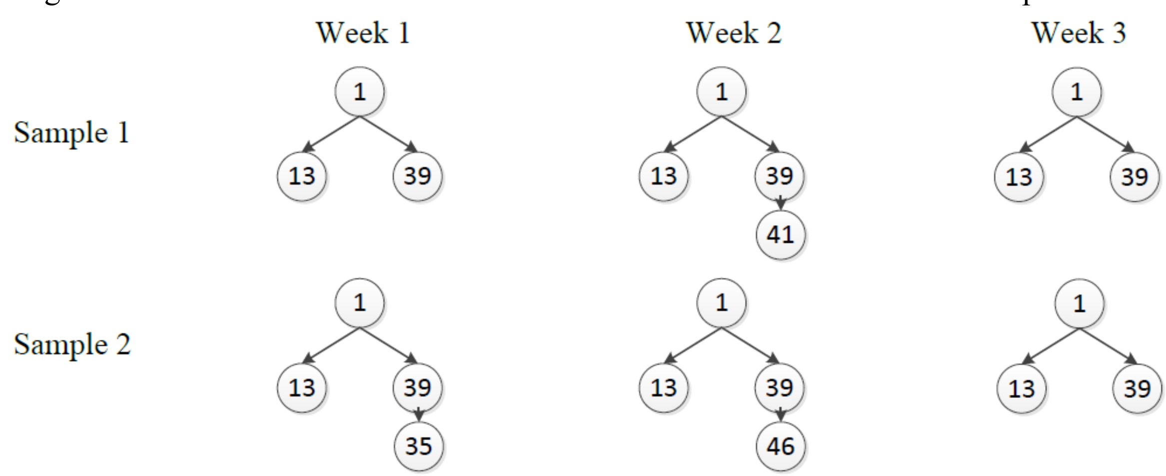

Figure 5.5. PAGs for Two Sets of Individuals and Three Time Periods with Alpha = 0.05

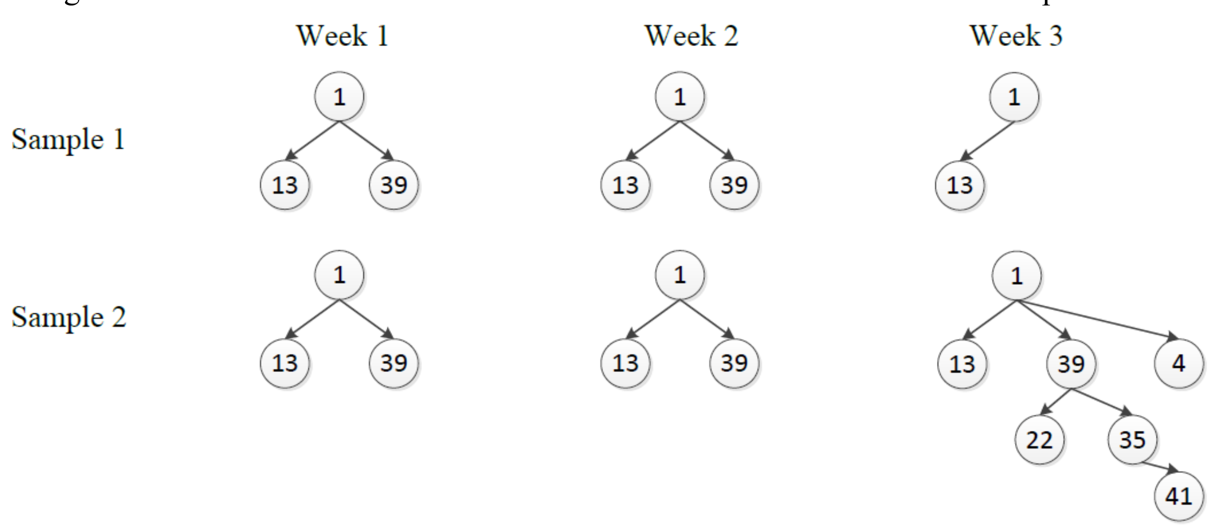

The PAGs are drawn based on conditional independence tests with a threshold for significance fixed at a certain level (alpha) to control type-1 errors in the statistical hypothesis testing framework. The level of alpha could be regarded as a trade-off between the probability of having an error in independence and the power of detecting dependence. As a result, PAGs estimated on the same observation but with different levels of alpha might exhibit different patterns. As seen in Figure 5.5, to examine the impact of alpha, we relax the alpha from 0.01 to 0.05 and redraw PAGs in the same way as in Figure 5.4.

Figure 5.6. PAG of Top-100 Popular Apps

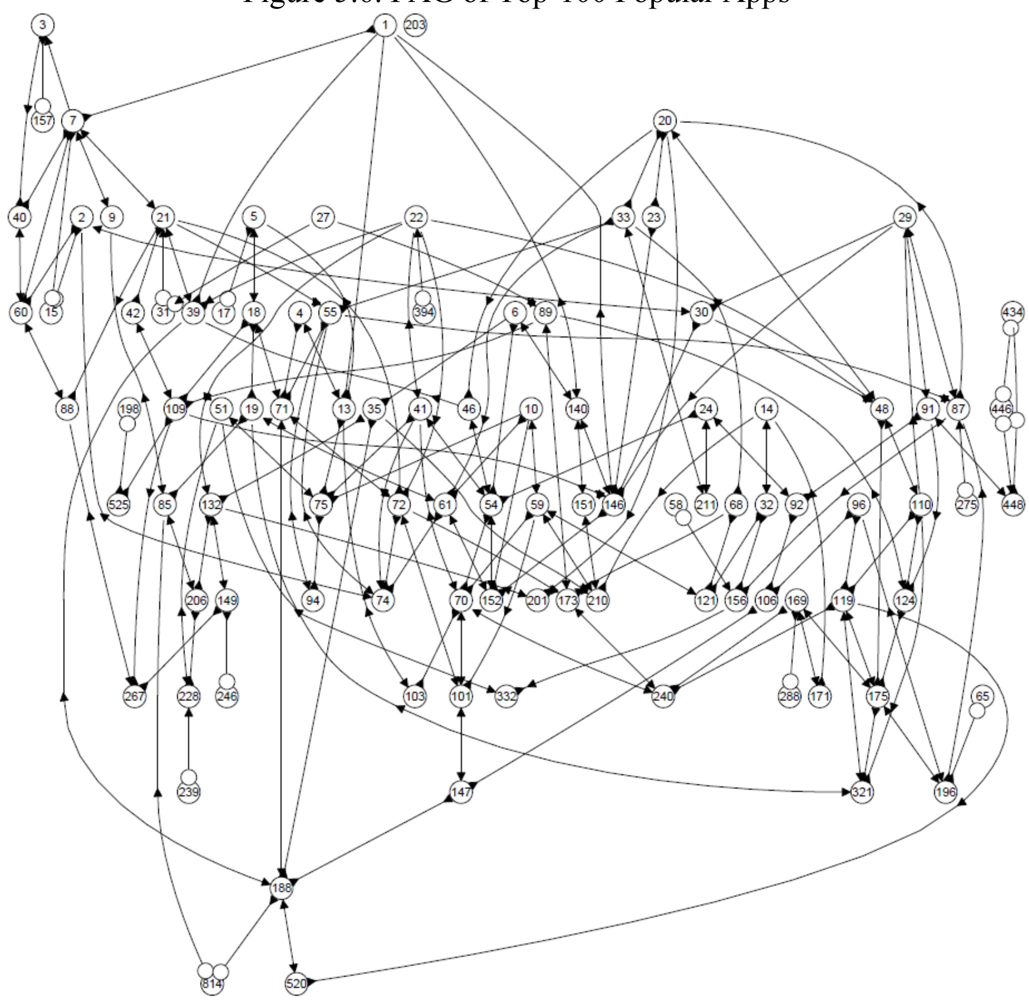

This figure also shows a high level of consistency of the causal structure. The majority of the graphs show a direct causal effect on Taobao (App 13) and Tencent News (App 39). The graph of Sample 1 in Week 3 does not have a causal path on Tencent News. Instead, it exhibits a bi-directed edge between WeChat and Tencent News. In addition, the graph of Sample 2 in Week 3 shows additional causal paths, including a direct causal effect on App 4. This is not surprising, however, because, as we increase the level alpha, we will have more vertices connected because the power of detecting the dependency signal increases.

The third robustness test is for the set of apps that we use for estimation. Additional information of app usage would provide more information on causal relationship identification. Therefore, robust causal relationships should stay constant if we increase the size of the vertices set. Specifically, we estimate two more PAGs with the top-100 frequently used apps and top-300 frequently used apps, correspondingly shown in Figure 5.6 and Figure 5.7, with the alpha as fixed at 0.01. The estimation for the PAG with the top-300 apps is implemented with the RFCI algorithm due to the infeasibility of applying the FCI to high dimensional data.

Figure 5.7. PAG of Top-300 Popular Apps

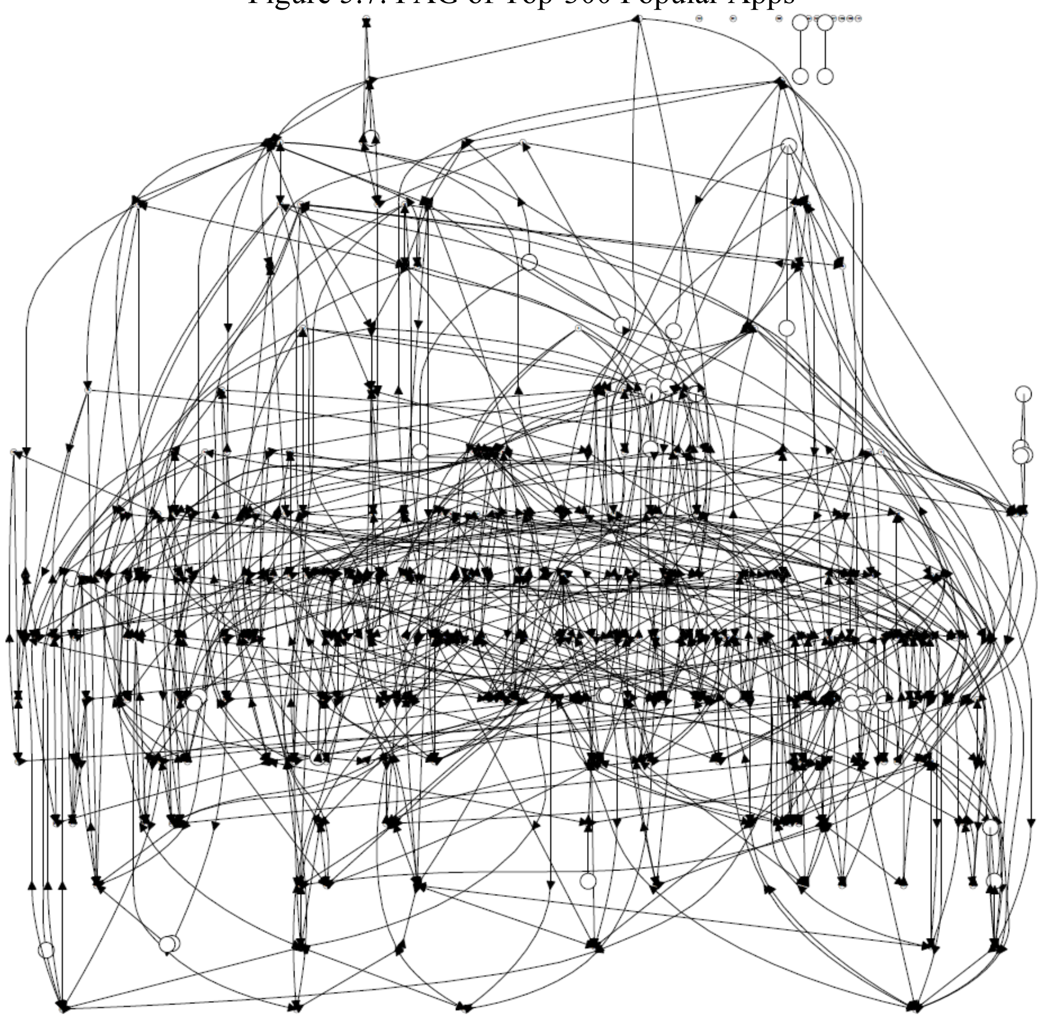

Due to the large scale of the vertices, the readability of the graph can be difficult. We examine the adjacent matrix and find the existence of causal paths from both WeChat to Taobao and to Tencent News, as seen in Figure 5.4 and Figure 5.5. Specifically, the PAG of the top-100 apps shows causal paths from WeChat to Taobao and to Tencent News as the only causal paths, which is exactly the same as seen in the PAGs of the top-50 apps. Further the PAG of the top-300 apps has causal paths to Taobao and to Tencent News as the only two direct causal paths. These consistencies suggest that our original model for the top-50 apps is able to capture most of spillover effects of WeChat. The PAG of the top-300 apps, however, has additional indirect causal paths to two apps, one of which is not included in the PAG of either the top-50 apps or of the top-100 apps. However, we reserve a conservative attitude toward these two causal paths for the following two reasons: (1) For the usage distribution of those less popular apps (of the top-100 popular apps set), it might be difficult to approximate the Gaussian distribution even after logarithm transformation. As we noted, when we took the logarithm of the app usage, if there was a great deal of zero usage, it could cause enormous skewness; and (2) RFCI-PAG is recognized as a super-graph of FCI and has weaker meaning in regard to the presence of edges than does the FCI, as shown in Colombo et al. (2012). Both reasons cast doubt on the robustness of these two causal effects.

5.5.2 *Check Quantitative Results*

As discussed in this section, we conduct a robustness check for the scale of causal effects. Specifically, we estimate causal effects from the data of distinct samples and time periods. Note that the estimation is based on the learned structure in the graphical results, and given a PAG, the specification for learning the graph has no impact on the quantitative estimation results. Therefore, there is no need to investigate the robustness of the alpha level or size of the vertices.

We first estimate spillover effects of WeChat on Tencent News and Taobao with distinct samples across different time periods separately, using the main model (5). The estimation results are shown in Table 5.18. Our results suggest a high degree of consistency across distinct samples and time periods. In all specifications of samples, the spillover effects on both Tencent News and Taobao are estimated to be positive, with the effect on Taobao as stronger quantitatively. The scales of effects are quite close among all six samples. The consistency of results based on different samples proves the robustness of our quantitative estimation.

Table 5.18. Comparing Quantitative Results Separate Samples

		Week1		Week2		Week3	
		Tencent News (13)	Taobao (39)	Tencent News (13)	Taobao (39)	Tencent News (13)	Taobao (39)
Sample 1	β_1	0.35*** (0.02)	0.40*** (0.03)	0.32*** (0.12)	0.37*** (0.02)	0.32*** (0.02)	0.33*** (0.03)
	c	0.20* (0.10)	0.35* (0.14)	0.12 (0.08)	0.09 (0.11)	0.18* (0.08)	0.24* (0.12)
Sample 2	β_1	0.38*** (0.02)	0.39*** (0.03)	0.36*** (0.02)	0.33*** (0.03)	0.35*** (0.02)	0.35*** (0.03)
	c	0.25* (0.10)	0.35* (0.15)	0.09 (0.09)	0.23 (0.12)	0.14 (0.08)	0.28* (0.11)

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05

Finally, note that pooling those six samples generates a sample of 1,200 individual smartphone users with repeated measures longitudinally. This pooled sample provides us with the opportunity to tease out individual-specific factors and time-specific factors to alleviate confounding bias. Note that our model suggests that no confounder exists on the causal paths from WeChat to Tencent News and to Taobao. Therefore, we expect estimates of parameters in a model with controlled individual-specific factors and time-specific factors to be similar to the estimates in former specifications. We control individual-specific factors and time-specific factors by adding fixed effects and specify the model as follows:

$$y_{it} = \beta_1' X_{it} + \beta_2' Z_{it} + c + \xi_i + \tau_t + \varepsilon_{it} \tag{5.55}$$

where ε_{it} are unobserved error terms following a Gaussian distribution. ξ_i and τ_t capture individual-specific unobserved effects and time-specific unobserved effects, respectively. Z_{it} is an empty set based on the GAC and GBC when estimating causal effects of WeChat on Tencent News and on Taobao. In addition, we estimate the causal effects by applying an OLS model without fixed effects on pooled data for comparison. We report the estimates in Table 5.19.

As we expect, parameter estimates for causal effects in Model (6) are very close to those of the original model (5). This implies the non-existence of a confounder that is encoded in the graphical model and further supports the robustness of our quantitative results.

Table 5.19. Spillover Effects Based on Pooled Sample

	Parameter	Tencent News (13)		Taobao (39)	
		FE	Pooled OLS	FE	Pooled OLS
Association	β_1	0.33***	0.35***	0.32***	0.37***
		(0.01)	(0.01)	(0.02)	(0.01)
	c	-0.48*	0.15***	-1.05	0.25***
		(0.44)	(0.04)	(0.61)	(0.05)
	FE	Not Report	NA	Not Report	NA
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05					

5.6 CONCLUSIONS, LIMITATIONS, AND FUTURE RESEARCH

The instant messaging app WeChat exhibits a mega status in the app market and exhibits dominance in terms of usage among Chinese smartphone users. However, its externality toward other types of apps has not received sufficient attention. Research is needed to investigate the spillover effects of such apps and to determine implications for the value that it creates for its developers as well as the value that it delivers to developers of other apps.

We combine a state-of-the-art machine learning method with an econometric approach to study the spillover effects of WeChat. Specifically, we apply an FCI-PAG method to determine the causal structure of app usage from observational data and estimate the spillover effects quantitatively based on the graphical outputs, the generalized back-door criterion, and generalized adjustment criterion. By applying our model to the app usage data of 600 Chinese smartphone users, we identify the set of apps that causally receive spillover effects from WeChat, the set of apps that shows association with WeChat due to observed or unobserved confounders, and the set of apps whose usage are independent of that of WeChat. We find that, counterintuitive to the belief of the industry, WeChat has quite limited external effects on the usage of other apps: among the top-50 and the top-100 apps, only two, Tencent News and Taobao, are shown to be causally positively affected by the usage of WeChat. Even when we extend the set to 300 apps, only these two apps receive spillover effects directly. The rest receive no causal effects from WeChat. To illustrate the importance of determining the causal structure and the value of quantitative information encoded in a graphical model, we further intentionally specify the econometric model with an incorrect adjustment set to show the erroneous estimation without the graphical results in the first stage.

Finally, we present the robustness of this approach by conducting a comparison of graphical estimates and quantitative estimates across samples in different time periods with different individuals. Using a pooled sample with repeated measure of individuals, we estimate the model with individual- and time-specific effects controlled to show the robustness of visible edges. In sum, this empirical study is the first to examine spillover effects of a mega app, such as WeChat. It provides researchers and app developers with a causal understanding, which is deeper than that provided with a superficial association explanation, and contributes to the analysis of attribution and decision-making about collaboration.

This paper is also the first to apply recent developments in machine learning-enabled causal inference models, such as FCI-PAG, plus a GAC and/or GBC estimation approach in business and economic research. Compared with past research methods, our approach relaxes the need for assumptions to identify causal effects with observational data but incurs a cost for obtaining additional information (hidden variables) when determining the causal structure. However, because data have become increasingly less expensive in this age of big data, this approach has the potential to be widely applied in estimating causal effects in business analytics research. Our work, as pioneering research that applies FCI-PAG plus GAC and/or GBC estimation, not only presents the spillover effects of WeChat but also shows a good fit of this advanced method in the context of business analytics research.

Our research is subject to limitations. These limitations typically relate to restrictions of the integration of the FCI-PAG-GAC or GBC approach and econometric methods, which, in turn, opens up avenues for future business analytics research. First, a more flexible model might be developed to allow for non-Gaussian distributed data. In our context, we use a log transformation to approximate our data to Gaussian. More complicated cases, such as ordinal choice data, however, might require a nonparametric graphical model to estimate. Second, even when we estimate a fixed-effects model with individual- and time-specific factors separately in the robustness check section, we notice that such factors cannot be added into the graphical estimation (first-stage estimation) due to the challenge of the independence test between Gaussian-distributed variables and dummy variables. A more generalized model that allows fixed effects might need to be developed to provide a more consistent (between graphical and quantitative models) and accurate estimation with repeated measurements graphically. Third, more graphical model-based econometric tools should be developed to help to estimate or validate the outputs from an FCI-

PAG-GAC/GBC approach. For example, a searching algorithm for generalized conditional instrument variables that works in a DAG/CPDAG setting should be extended to an MAG/PAG setting to help to validate the visible edge.

These three recommendations, based on the limitations of our research, would help to further develop the connection between econometric and graphical models. Given the similarities of the nature of these two methods, more complete integration should be promising in future research. Further, our method can be easily adapted to other research contexts, such as online social networks and recommendation systems. Given the power of drawing causal inferences from observational data, we expect that more fruitful applications of this approach would contribute to a better analytical understanding of business.

BIBLIOGRAPHY

[1]. Arcidiacono, P., 2005. Affirmative action in higher education: How do admission and financial aid rules affect future earnings?. Econometrica, 73(5), pp.1477-1524.

[2]. Ba, S., Pavlou, P.A. 2002. Evidence of the effect of trust building technology in electronic markets: Price premiums and buyer behavior. MIS quarterly, pp.243-268.

[3]. Bajari, P., Benkard, C.L. and Levin, J., 2007. Estimating dynamic models of imperfect competition. Econometrica, 75(5), pp.1331-1370.

[4]. Bapna, R., Jank, W., Shmueli, G. 2008. Consumer surplus in online auctions. Information Systems Research, 19(4), pp.400-416.

[5]. Berry, S., Levinsohn, J., Pakes, A. 1995. Automobile prices in market equilibrium. Econometrica: Journal of the Econometric Society, pp.841-890.

[6]. Berry, S., Levinsohn, J., Pakes, A. 1998. Differentiated products demand systems from a combination of micro and macro data: The new car market (No. w6481). National bureau of economic research.

[7]. Berry, S., Levinsohn, J., Pakes, A. 1999. Voluntary export restraints on automobiles: Evaluating a trade policy. American Economic Review, pp.400-430.

[8]. Bockstedt, J., Goh, K.H. 2011. Seller strategies for differentiation in highly competitive online auction markets. Journal of Management Information Systems, 28(3), pp.235-268.

[9]. Bolton, G., Loebbecke, C., Ockenfels, A. 2008. Does competition promote trust and trustworthiness in online trading? An experimental study. Journal of Management Information Systems, 25(2), pp.145-170.

[10]. Brynjolfsson, E., Hu, Y., Smith, M.D. 2003. Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. Management Science, 49(11), pp.1580-1596.

[11]. Buntine, W., 1991, July. Theory refinement on Bayesian networks. In Proceedings of the Seventh conference on Uncertainty in Artificial Intelligence (pp. 52-60). Morgan Kaufmann Publishers Inc.

[12]. California Public Utilities Commission, 2013. Decision adopting rules and regulations to protect public safety while allowing new entrants to the transportation industry. Web, http://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M077/K112/77112285.PDF, 9 Sept 2013.

[13]. Campos, L.M.D., Gámez Martín, J.A. and Puerta Castellón, J.M., 2002. Learning Bayesian networks by ant colony optimisation: searching in two different spaces. Mathware & soft computing. 2002 Vol. 9 Núm. 2 [-3].

[14]. Carare, O., 2012. The impact of bestseller rank on demand: Evidence from the app market. International Economic Review, 53(3), pp.717-742.

[15]. Chan, C. (2015, August 6). When One App rules them all: The case of WeChat and mobile in china. Retrieved August 15, 2016, from A16Z, http://a16z.com/2015/08/06/wechat-china-mobile-first/

[16]. Chan, Tat Y., and Barton H. Hamilton. "Learning, private information, and the economic evaluation of randomized experiments." Journal of Political Economy 114.6 (2006): 997-1040.

[17]. Chen, T., Sun, B. and Singh, V., 2009. An empirical investigation of the dynamic effect of Marlboro's permanent pricing shift. Marketing Science, 28(4), pp.740-758.

[18]. Colombo, D., Maathuis, M.H., Kalisch, M. and Richardson, T.S., 2012. Learning high-dimensional directed acyclic graphs with latent and selection variables. The Annals of Statistics, pp.294-321.

[19]. Cooper, G.F. and Herskovits, E., 1992. A Bayesian method for the induction of probabilistic networks from data. Machine learning, 9(4), pp.309-347.

[20]. Cormack, M. (2015, February 10). WeChat's impact: A report on WeChat platform data. Retrieved August 15, 2016, from technode, http://technode.com/2015/02/10/wechat-impact-report/

[21]. Cowie, J., Oteniya, L. and Coles, R., 2007. Particle swarm optimisation for learning Bayesian networks. In ICCIIS 2007, World Congress on Engineering, WCE 2007 (pp. 71-76). Newswood Limited/International Association of Engineers (IAENG).

[22]. Crawford, G.S. and Shum, M., 2005. Uncertainty and learning in pharmaceutical demand. Econometrica, 73(4), pp.1137-1173.

[23]. Davis, P. 2006. Spatial competition in retail markets: movie theaters. The RAND Journal of Economics, 37(4), pp.964-982.

[24]. De Campos, L.M. and Huete, J.F., 2000. A new approach for learning belief networks using independence criteria. International Journal of Approximate Reasoning, 24(1), pp.11-37.

[25]. Dimoka, A., Hong, Y., Pavlou, P.A. 2012. On product uncertainty in online markets: Theory and evidence. MIS quarterly, 36.

[26]. Dubé, J.P.H., Fox, J.T., Su, C.L. 2009. Improving the numerical performance of blp static and dynamic discrete choice random coefficients demand estimation (No. w14991). National Bureau of Economic Research.

[27]. Elwert, F., 2013. Graphical causal models. In Handbook of causal analysis for social research (pp. 245-273). Springer Netherlands.

[28]. Erdem, T. and Keane, M.P., 1996. Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. Marketing science, 15(1), pp.1-20.

[29]. Erdem, T. and Sun, B., 2002. An empirical investigation of the spillover effects of advertising and sales promotions in umbrella branding. Journal of Marketing Research, 39(4), pp.408-420.

[30]. Erdem, T., Keane, M.P. and Sun, B., 2008. A dynamic model of brand choice when price and advertising signal product quality. Marketing Science, 27(6), pp.1111-1125.

[31]. expands-its-on-demand-home-services-marketplace-adds-support-for-custom-service-requests.

[32]. Falaki, H., Lymberopoulos, D., Mahajan, R., Kandula, S. and Estrin, D., 2010, November. A first look at traffic on smartphones. In Proceedings of the 10th ACM SIGCOMM conference on Internet measurement (pp. 281-287). ACM.

[33]. Garg, R. and Telang, R., 2012. Inferring app demand from publicly available data. MIS Quarterly, Forthcoming.

[34]. Ghose, A. 2009. Internet exchanges for used goods: An empirical analysis of trade patterns and adverse selection. MIS Quarterly, pp.263-291.

[35]. Ghose, A. and Han, S.P., 2011. An empirical analysis of user content generation and usage behavior on the mobile Internet. Management Science, 57(9), pp.1671-1691.

[36]. Ghose, A. and Han, S.P., 2014. Estimating demand for mobile applications in the new economy. Management Science, 60(6), pp.1470-1488.

[37]. Ghose, A. and Ipeirotis, P.G., 2011. Estimating the helpfulness and economic impact of product reviews: Mining text and reviewer characteristics. IEEE Transactions on Knowledge and Data Engineering, 23(10), pp.1498-1512.

[38]. Ghose, A., Han, S.P. 2011. An empirical analysis of user content generation and usage behavior on the mobile Internet. Management Science, 57(9), pp.1671-1691.

[39]. Ghose, A., Han, S.P. 2014. Estimating demand for mobile applications in the new economy. Management Science, 60(6), pp.1470-1488.

[40]. Ghose, A., Ipeirotis, P.G., Li, B. 2012. Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. Marketing Science, 31(3), pp.493-520.

[41]. Granados, N., Gupta, A., Kauffman, R.J. 2012. Online and offline demand and price elasticities: Evidence from the air travel industry.Information Systems Research, 23(1), pp.164-181.

[42]. Greene, W.H. and Hensher, D.A., 2003. A latent class model for discrete choice analysis: contrasts with mixed logit. Transportation Research Part B: Methodological, 37(8), pp.681-698.

[43]. Gregg, D.G., Walczak, S. 2008. Dressing your online auction business for success: An experiment comparing two eBay businesses. MIS Quarterly, pp.653-670.

[44]. Gupta, S. and Chintagunta, P.K., 1994. On using demographic variables to determine segment membership in logit mixture models. Journal of Marketing Research, pp.128-136.

[45]. Heckerman, D., Geiger, D. and Chickering, D.M., 1995. Learning Bayesian networks: The combination of knowledge and statistical data. Machine learning, 20(3), pp.197-243.

[46]. Hennig-Thurau, T., Gwinner, K.P., Walsh, G. and Gremler, D.D., 2004. Electronic word-of-mouth via consumer-opinion platforms: What motivates consumers to articulate themselves on the Internet?. Journal of interactive marketing, 18(1), pp.38-52.

[47]. Hensher, D.A. and Stopher, P.R. eds., 1979. Behavioural travel modelling. Taylor & Francis.

[48]. Ho, Y.-C., Wu, J. and Tan, Y., 2014. Disconfirmation Effect on Online Rating Behavior: A Dynamic Analysis. University of Washington Working Paper.

[49]. Ho, Y.C.C., Wu, J. and Tan, Y., 2016. A Structural Model of Disconfirmation Effect on Online Rating Behavior. Information Systems Research, forthcoming.

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

[51]. Huang, Y. and Valtorta, M., 2006, July. Identifiability in causal Bayesian networks: A sound and complete algorithm. In PROCEEDINGS OF THE NATIONAL CONFERENCE ON ARTIFICIAL INTELLIGENCE (Vol. 21, No. 2, p. 1149). Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999.

[52]. Huang, Y., Vir Singh, P. and Srinivasan, K., 2014. Crowdsourcing new product ideas under consumer learning. Management science, 60(9), pp.2138-2159.

[53]. Juggernaut. 2015. How to Succeed with a Services Marketplace: Choose the Model That Fits You Right! Juggernaut., 09 Nov. 2015. Accessed 03 May 2016, http://nextjuggernaut.com/blog/how-to-succeed-with-a-services-marketplace-choose-the-model-that-fits-you-right/.

[54]. Kabli, R., Herrmann, F. and McCall, J., 2007, July. A chain-model genetic algorithm for Bayesian network structure learning. In Proceedings of the 9th annual conference on Genetic and evolutionary computation (pp. 1264-1271). ACM.

[55]. Kalisch, M. and Bühlmann, P., 2007. Estimating high-dimensional directed acyclic graphs with the PC-algorithm. Journal of Machine Learning Research,8(Mar), pp.613-636.

[56]. Kalisch, M., Mächler, M., Colombo, D., Maathuis, M.H. and Bühlmann, P., 2012. Causal inference using graphical models with the R package pcalg. Journal of Statistical Software, 47(11), pp.1-26.

[57]. Koller, D. and Friedman, N., 2009. Probabilistic graphical models: principles and techniques. MIT press.

[58]. Larrañaga, P., Murga, R., Poza, M. and Kuijpers, C., 1996. Structure learning of Bayesian networks by hybrid genetic algorithms. In Learning from Data (pp. 165-174). Springer New York.

[59]. Lee, D., Hosanagar, K., and Nair, H. 2013. The effect of advertising content on consumer engagement: Evidence from Facebook. Available at SSRN 2290802.

[60]. Lee, Y.J., Hosanagar, K. and Tan, Y., 2015. Do I follow my friends or the crowd? Information cascades in online movie ratings. Management Science, 61(9), pp.2241-2258.

[61]. Lerner, J. and Tirole, J., 2002. Some simple economics of open source. The Journal of Industrial Economics, 50(2), pp.197-234.

[62]. Leslie, P. 2004. Price discrimination in Broadway theater. RAND Journal of Economics, pp.520-541.

[63]. Li, C. and Zhao, J., 2015. Humanizing travel: how e-hail apps transform stakeholder relationships in the US taxi industry. In Transportation Research Board 94th Annual Meeting (No. 15-0680).

[64]. Li, S., Srinivasan, K., Sun, B. 2009. Internet auction features as quality signals. Journal of Marketing, 73(1), pp.75-92.

[65]. Li, X. and Hitt, L.M., 2008. Self-selection and information role of online product reviews. Information Systems Research, 19(4), pp.456-474.

[66]. Li, Y., Xia, T. and Duan, H., 2014. The impact on taxi industry of taxi-calling mobile apps in shanghai. In Transportation Research Board 93rd Annual Meeting (No. 14-3867).

[67]. Maathuis, M.H. and Colombo, D., 2015. A generalized back-door criterion. The Annals of Statistics, 43(3), pp.1060-1088.

[68]. McCullagh, P., 1980. Regression models for ordinal data. Journal of the Royal Statistical Society. Series B (Methodological), pp.109-142.

[69]. McFadden, D., 1973. Conditional logit analysis of qualitative choice behavior.

[70]. Moe, W.W. and Schweidel, D.A., 2012. Online product opinions: Incidence, evaluation, and evolution. Marketing Science, 31(3), pp.372-386.

[71]. Moe, W.W. and Trusov, M., 2011. The value of social dynamics in online product ratings forums. Journal of Marketing Research, 48(3), pp.444-456.

[72]. Mudambi, S.M. and Schuff, D., 2010. What makes a helpful review? A study of customer reviews on Amazon. com. MIS Quarterly, 34(1), pp.185-200.

[73]. Neapolitan, R.E., 2004. Learning Bayesian networks.

[74]. Nevo, A. 2000. A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand. Journal of Economics & Management Strategy,9(4), pp.513-548.

[75]. Nevo, A. 2001. Measuring market power in the ready-to-eat cereal industry. Econometrica, 69(2), pp.307-342.

[76]. Osawa, J., 2015. Why China Taxi Apps Are Talking Merger. Web, Retrieved March 26, 2015, from http://blogs.wsj.com/digits/2015/02/13/why-china-taxi-apps-are-talking-merger/

[77]. Pan, Y. and Zhang, J.Q., 2011. Born unequal: a study of the helpfulness of user-generated product reviews. Journal of Retailing, 87(4), pp.598-612.

[78]. Pavlou, P.A., Dimoka, A. 2006. The nature and role of feedback text comments in online marketplaces: Implications for trust building, price premiums, and seller differentiation. Information Systems Research, 17(4), pp.392-414.

[79]. Pearl, J., 1982, August. Reverend Bayes on inference engines: A distributed hierarchical approach. In AAAI (pp. 133-136).

[80]. Pearl, J., 1993. [Bayesian Analysis in Expert Systems]: Comment: Graphical Models, Causality and Intervention. Statistical Science, 8(3), pp.266-269.

[81]. Pearl, J., 1995. Causal diagrams for empirical research. Biometrika, 82(4), pp.669-688.

[82]. Pearl, J., 2011. Bayesian networks. Department of Statistics, UCLA,

[83]. Pearl, T.V.J., 1991. Equivalence and synthesis of causal models. In Proceedings of Sixth Conference on Uncertainty in Artificial Intelligence (pp. 220-227).

[84]. Peng, J., Wang, P., Zhou, N. and Zhu, J., 2012. Partial correlation estimation by joint sparse regression models. Journal of the American Statistical Association.

[85]. Perez, S. 2015. Amazon Expands Its On-Demand Home Services Marketplace, Adds Support For Custom Service Requests. TechCrunch. Accessed 03 May 2016, http://techcrunch.com/2015/07/22/amazon-

[86]. Perez, S., Etherington, D. 2015. Amazon's On-Demand Services Marketplace Launches Monday. TechCrunch. Accessed 03 May 2016, http://techcrunch.com/2015/03/25/amazons-on-demand-services-marketplace-launches-monday.

[87]. Perković, E., Textor, J., Kalisch, M. and Maathuis, M.H., 2015. A complete generalized adjustment criterion. arXiv preprint arXiv:1507.01524.

[88]. Petrin, A. 2001. Quantifying the benefits of new products: The case of the minivan (No. w8227). National Bureau of Economic Research.

[89]. Ramsey, J., Zhang, J. and Spirtes, P.L., 2012. Adjacency-faithfulness and conservative causal inference. arXiv preprint arXiv:1206.6843.

[90]. Rayle, L., Shaheen, S., Chan, N., Dai, D. and Cervero, R., 2014. App-Based, On-Demand Ride Services: Comparing Taxi and Ridesourcing Trips and User Characteristics in San Francisco University of California Transportation Center (UCTC). UCTC-FR-2014-08.

[91]. Richardson, T. and Spirtes, P., 2002. Ancestral graph Markov models. Annals of Statistics, pp.962-1030.

[92]. Rust, J., 1987. Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. Econometrica: Journal of the Econometric Society, pp.999-1033.

[93]. Schlosser, A.E., 2005. Posting versus lurking: Communicating in a multiple audience context. Journal of Consumer Research, 32(2), pp.260-265.

[94]. Schlosser, A.E., White, T.B., Lloyd, S.M. 2006. Converting web site visitors into buyers: how web site investment increases consumer trusting beliefs and online purchase intentions. Journal of Marketing, 70(2), pp.133-148.

[95]. Schwarz, G., 1978. Estimating the dimension of a model. The annals of statistics, 6(2), pp.461-464.

[96]. Scutari, M. and Denis, J.B., 2014. Bayesian networks: with examples in R. CRC Press.

[97]. Seetharaman, D. 2014. Exclusive: Amazon.com plans local services marketplace this year - sources. Reuters. Accessed 03 May 2016, http://www.reuters.com/article/us-amazon-com-services-idUSKBN0EL20S20140610.

[98]. Sen, S. and Lerman, D., 2007. Why are you telling me this? An examination into negative consumer reviews on the web. Journal of Interactive Marketing, 21(4), pp.76-94.

[99]. Shachar, R., Anand, B.N. 1998. The effectiveness and targeting of television advertising. Journal of Economics & Management Strategy, 7(3), pp.363-396.

[100]. Shen, W., Hu, Y.J. and Rees, J., 2015. Competing for attention: An empirical study of online reviewers' strategic behaviors. MIS Quarterly, 39(3), pp.683-696.

[101]. Shpitser, I. and Pearl, J., 2008. Complete identification methods for the causal hierarchy. Journal of Machine Learning Research, 9(Sep), pp.1941-1979.

[102]. Sobel, M.E., 1996. An introduction to causal inference. Sociological Methods & Research, 24(3), pp.353-379.

[103]. Sohu IT. 2015. New Market Structure for Domestic Mobile Phone: Increment in 4G Phone, Decrement in Contract Bundled Phone. Accessed 03 May 2016, http://it.sohu.com/20141017/n405198791.shtml

[104]. Spirtes, Peter, Clark N. Glymour, and Richard Scheines. Causation, prediction, and search. MIT press, 2000.

[105]. Sugden, J., Aditi M. and Douglas M., 2014. Uber Under Attack Around the Globe. The Wall Street Journal, Web, Retrieved March 26, 2015, from http://www.wsj.com/articles/india-advises-all-states-to-ban-uber-and-other-car-hailing-services-1418119647.

[106]. Tan, F.T.C., Tan, B. and Land, L., 2015, January. The affordance of gamification in enabling a digital disruptor: a case study of the goCatch taxi booking app. In System Sciences (HICSS), 2015 48th Hawaii International Conference on (pp. 1197-1206). IEEE.

[107]. Tang, Z., Hu, Y., Smith, M.D. 2008. Gaining trust through online privacy protection: Self-regulation, mandatory standards, or caveat emptor. Journal of Management Information Systems, 24(4), pp.153-173.

[108]. Tangel, A., 2014. Ride-Sharing Service Lyft Defies New York City, State Regulators. The Wall Street Journal, Web, Retrieved March 26, 2015, from http://www.wsj.com/articles/ride-sharing-service-lyft-defies-new-york-city-state-regulators-1405042425.

[109]. Tian, J. and Pearl, J., 2002, August. A general identification condition for causal effects. In AAAI/IAAI (pp. 567-573).

[110]. Tongaonkar, A., Dai, S., Nucci, A. and Song, D., 2013, March. Understanding mobile app usage patterns using in-app advertisements. In International Conference on Passive and Active Network Measurement (pp. 63-72). Springer Berlin Heidelberg.

[111]. Tsamardinos, I., Brown, L.E. and Aliferis, C.F., 2006. The max-min hill-climbing Bayesian network structure learning algorithm. Machine learning,65(1), pp.31-78.

[112]. Wang, T., Touchman, J.W. and Xue, G., 2004, August. Applying two-level simulated annealing on Bayesian structure learning to infer genetic networks. In Computational Systems Bioinformatics Conference, 2004. CSB 2004. Proceedings. 2004 IEEE (pp. 647-648). IEEE.

[113]. Wells, J.D., Valacich, J.S., Hess, T.J. 2011. What Signals Are You Sending? How Website Quality Influences Perceptions of Product Quality and Purchase Intentions. MIS quarterly, 35(2), pp.373-396.

[114]. Wu, C., Hai C., Tat Y. C. and Xianghua L., 2015. The economic value of online reviews. Marketing Science, forthcoming.

[115]. Xu, Y., Lin, M., Lu, H., Cardone, G., Lane, N., Chen, Z., Campbell, A. and Choudhury, T., 2013, September. Preference, context and communities: a multi-faceted approach to predicting smartphone app usage patterns. In Proceedings of the 2013 International Symposium on Wearable Computers (pp. 69-76). ACM.

[116]. Yang, M., Ren, Y., and Adomavicious, G., 2016 Understanding word-of-mouth and customer engagement on Facebook business pages. University of Minnesota Working Paper.

[117]. Yin, D., Mitra, S. and Zhang, H., 2016. Research Note—When Do Consumers Value Positive vs. Negative Reviews? An Empirical Investigation of Confirmation Bias in Online Word of Mouth. Information Systems Research, 27(1), pp.131-144.

[118]. Ying, Y., Feinberg, F. and Wedel, M., 2006. Leveraging missing ratings to improve online recommendation systems. Journal of Marketing Research, 43(3), pp.355-365.

[119]. Yoganarasimhan, H. 2013. The value of reputation in an online freelance marketplace. Marketing Science, 32(6), pp.860-891.

[120]. Zhang, J., 2008. Causal reasoning with ancestral graphs. Journal of Machine Learning Research, 9(Jul), pp.1437-1474.

[121]. Zhang, J., 2008. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. Artificial Intelligence, 172(16), pp.1873-1896.

[122]. Zhang, J.Q., Craciun, G. and Shin, D., 2010. When does electronic word-of-mouth matter? A study of consumer product reviews. Journal of Business Research, 63(12), pp.1336-1341.

[123]. Zhao, Y., Yang, S., Narayan, V. and Zhao, Y., 2013. Modeling consumer learning from online product reviews. Marketing Science, 32(1), pp.153-169.

APPENDIX A

A1. Sieve Model

Given that the number of RIRs is a count variable and shown with severe over-dispersion and inflation at zero, we specify a zero-inflated negative binomial model to form the structure. The specification of the model is shown as follows:

$$\Pr\left(p_{irt}=y\right)=\left(\frac{\mu p_{irt}}{1+\mu p_{irt}}\right)\exp\left(y\ln\left(\frac{\mu p_{irt}}{\theta+\mu p_{irt}}\right)-\theta\ln\left(1+\frac{\mu p_{irt}}{\theta}\right)+\ln\left(\frac{\Gamma(y+\theta)}{\Gamma(y+1)\Gamma(\theta)}\right)\right)$$

$$\text{for } y>0, \tag{A1}$$

$$\Pr\left(p_{irt}=y\right)=\frac{1}{1+\mu p_{irt}}+\left(\frac{\mu p_{irt}}{1+\mu p_{irt}}\right)\exp\left(-\theta\ln\left(1+\frac{\mu p_{irt}}{\theta}\right)\right) \text{ for } y=0, \tag{A2}$$

where p_{irt} is the expected number of positive RIRs for review i at time t from the date that the rating is posted; y is a natural number; μp_{irt} is the expected value parameter for positive RIR; and θ is the dispersion parameter for positive RIR.

There are multiple methods of specifying the expected value parameter μp_{irt} to allow flexibility. We use an approach to model this expectation by mimicking the idea of the logit sieve model, which uses a higher-order polynomial to allow more flexibility by approximating a Taylor expansion of the true parameter of μp_{irt}. Specifically, we model:

$$\mu p_{irt}=\exp\left(\gamma_p X_{irt}\right), \tag{A3}$$

where X_{irt} is a vector of the first to fourth orders and interactions of the following non-dummy variables and the first order of the dummy variables: (a) xt_{ti} duration from the date the rating is posted; (b) $I(xd_{ti})$ factor indicator of the day of the week of time t; (c) $I(pc_i)$ indicator of product type, with 1 as representing a film and 0 as representing a book; (d) pq_i product quality; (e) $xg_{ri} = (r_i - xq_i / 2)$ confirmation bias; (f) $\ln(pl_i + 1)$ log of one plus the number of existing ratings for the same product by the same type of content generator (Type 3 rating) before the rating is posted; (g) $\ln(pw_i + 1)$ log of one plus the number of cumulative watching/reading experiences related to the product by all types of users; (h) $\ln(pt_i + 1)$ log of one plus the cumulative number of watching/reading intentions related to the product reported by all types of users; (i) $\ln(pp_i + 1)$ log of one plus the number of product pictures posted by users; (j) $\ln(uc_i + 1)$ log of one plus the rater's incoming ties for rating i; (k) $\ln(uf_i + 1)$ log of one plus the self-reported number of films watched by the rater; (l) $\ln(ub_i + 1)$ log of one plus the self-reported number of books read by the rater; and (m) $\ln(ur_i + 1)$ log of one plus the cumulative number of reviews generated by the rater. Those variables include all of the variables relevant to RIR generation. In addition, for the purpose of consistency with the model-free finding shown in Section 4.3.2, we use confirmation bias instead of rating valence in this model.

We then form expectations in two stages. In the first stage, we estimate parameters for the above-referenced model by maximum likelihood and apply the same approach to recover the same set of parameters for an expected negative RIR for rating i at time t since posting d_{irt}, along with that of an expected mixed RIR for rating i at time t since posting n_{irt}. This stage mimics content generators' learning from past ratings and reviews. In the second stage, given the estimated parameters, we recover RIR expectations for the current rating and review by applying the estimated parameter with current period data to predict the corresponding RIR. This stage mimics the expectation formation conditional on the content generator's decision of a rating valence for the current stage. We test different orders of polynomials and find that the fourth order renders the best-fitting performance while maintaining a manageable parameter size

A2. Robustness Check

Estimates of alternative specifications are presented in Table A1. First, recall that we use a discount factor of 0.95 to fix the devaluation rate when raters calculate the present value of social capital. To check the potential impact of fixed discount factors, we replicate our model with a faster devaluation, with a discount factor of 0.80 (Model (5)) and our model, with a slower devaluation with a discount factor of 0.99 (Model (6)). Second, in the original model, we calculate expected social capital with a finite horizon of 80 days in the future, as 0.95^{80} is as small as 0.0165, which further discounts the RIR arrival to be very small. Such a horizon, however, might not be sufficient to approximate an infinite horizon, given that 0.0165 is not close enough to 0. We set the horizon to be double that of the original model as 160 days in Model (7), with weights for social capital arriving later and fewer than 0.95^{160} as 0.00028 to erase the concern that the horizon is not long enough. Finally, note that our model incorporates one predictive stage to estimate expected social capital gains. To validate the predictive results generated in Stage 1, we relocate 20% of the observations used in the parameter estimation (Stage 2) to Stage 1, which increases the size of the training sample by 20% and decreases the sample size of Stage 2 by 20%, thus estimating Model (8). Although this is not a typical predictive model validation, it is the best we can do to infer whether Stage 1 is robust and whether the estimation in Stage 1 will affect the estimation in Stage 2, given a limited sample size and the number of dimensions of the parameters.

The results from these alternative models show that not only the main qualitative features but also the quantitative estimates of our results are robust across our various specifications, except for some sign flipping related to insignificant estimates. This suggests that the limitation of the fixed discount factor, limited time horizon and predictive estimation does not significantly affect our final point estimate results.

Table A.1. Robustness Check

Variables	0.80 Discount Factor (5)	0.99 Discount Factor (6)	Longer Future Horizon (7)	Shorter Stage 2 Period (8)
Help	4.54***(1.43)	4.63** (1.58)	5.40*** (1.58)	3.65***(0.79)
Unhelp	-1.71 (2.16)	-2.30* (1.93)	-3.10* (2.09)	-2.41* (2.14)
Constant R1	11.51***(2.44)	11.40***(2.36)	10.26*** (2.96)	10.96** (4.26)
Constant R2	15.88***(3.05)	15.82***(2.54)	15.53*** (3.56)	16.50***(4.09)
Constant R3	18.41***(3.38)	18.42***(2.89)	18.85*** (3.99)	18.99***(4.11)
Constant R4	17.94***(3.14)	18.10***(2.76)	19.10*** (3.97)	17.30***(4.02)
Help^2	0.25* (0.20)	0.26* (0.25)	0.28* (0.19)	0.16* (0.14)
Unhelp^2	-0.84 (4.57)	-3.30 (3.97)	-1.89 (4.42)	0.03 (3.91)
Replies	4.69***(1.31)	5.03***(1.38)	5.85*** (1.81)	4.39***(0.73)
Replies^2	0.16 (0.28)	0.17 (0.30)	-0.19 (0.24)	-0.24* (0.22)
Constant T1	0.74***(0.03)	0.74***(0.03)	0.74*** (0.03)	0.74***(0.03)
Constant T2	0.78***(0.03)	0.78***(0.03)	0.78*** (0.03)	0.78***(0.04)
Constant T3	0.18* (0.10)	0.17* (0.09)	0.17** (0.08)	0.03 (0.07)
Constant T4	-4.15***(0.14)	-4.14***(0.14)	-4.14*** (0.12)	-4.07***(0.10)
Avg Quality	0.29***(0.09)	0.29***(0.09)	0.29*** (0.09)	0.25** (0.09)
Is Movie Review	0.29* (0.19)	0.29* (0.19)	0.29* (0.20)	0.28* (0.19)
Picture	0.06* (0.06)	0.07* (0.06)	0.06 (0.07)	0.03 (0.06)
Rating Volume	0.10 (0.23)	0.10 (0.23)	0.10 (0.23)	0.14 (0.26)
ShortRev Volume	0.08 (0.18)	0.08 (0.18)	0.09 (0.20)	0.12 (0.20)
LongRev Volume	0.12* (0.07)	0.12* (0.08)	0.11* (0.07)	0.17* (0.11)
Watched	-0.22* (0.19)	-0.22* (0.19)	-0.22* (0.19)	-0.26* (0.21)
Watching Intention	-0.03 (0.12)	-0.03 (0.13)	-0.03 (0.12)	-0.05 (0.10)
1star	-0.54***(0.06)	-0.54***(0.06)	-0.53*** (0.06)	-0.60***(0.08)
2star	-0.34***(0.05)	-0.34***(0.05)	-0.34*** (0.05)	-0.33***(0.05)
3star	-0.53***(0.05)	-0.54***(0.06)	-0.54*** (0.05)	-0.53***(0.05)
4star	-0.13** (0.05)	-0.13** (0.05)	-0.13** (0.05)	-0.12** (0.05)
Constant P	0.37 (0.69)	0.28 (0.70)	0.25 (1.03)	0.40 (1.59)
Outgoing Ties	0.38 (1.07)	0.45 (1.06)	0.57 (1.18)	0.76 (1.64)
Incoming Ties	0.53 (0.65)	0.44 (0.79)	0.48 (0.79)	0.28 (0.73)
Groups	5.63** (2.67)	6.14** (2.63)	6.12* (3.91)	9.25* (5.30)
Total Reviews	-5.32** (2.06)	-5.66** (2.08)	-5.63** (2.41)	-7.80* (4.34)
Outgoing Ties^2	-1.18* (0.69)	-1.30* (0.69)	-1.32* (0.77)	-2.07* (1.52)
Incoming Ties^2	13.33* (7.39)	14.87* (7.52)	14.93* (10.94)	25.43** (11.32)
Groups^2	5.48** (2.57)	5.96** (2.51)	6.00* (3.65)	9.01* (4.80)
Total Reviews^2	3.03** (1.32)	3.27** (1.33)	3.33** (1.53)	4.05* (2.37)
Sample Size	2940	2940	2940	2365
Horizon Length	80	80	180	80
Discount Factor	0.80	0.99	0.95	0.95

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

A3. Marginal Effect of Expected RIRs

Our results show that individuals' rating decisions are based not only on perceived quality but also on expected social capital. However, the analysis so far is inadequate to generate quantitative implications for the magnitude of expected RIR impact, e.g., the marginal effect of expected RIRs, due to the nonlinearity of the model. Therefore, our second set of simulations is set to measure the quantitative impact of RIRs by experimenting with how significantly the ratings distribution will be biased if individuals have a different expectation of receiving RIRs. Note that the RIR-generation process is not the focus of this research, and, thus, it is modeled in a flexible manner. Therefore, our simulation focuses exclusively on the expected RIR impact on rating, not on how to adjust raters' expectations of RIRs.

To test the impact of adding RIRs, we conduct four simulation experiments. First, we note that adding the expectation of "helpful" for a targeted level of ratings increases raters' expected social capital conditional on that level of rating, thus encouraging raters to rate at that level. Assuming a goal of increasing the rating valence of the platform, we add 10 more marginal expected "helpfuls" under the conditions of 4- and 5-star ratings so that we can quantify the marginal effect on ratings of increasing expected social capital for a targeted level of ratings. Second, in addition to increasing expected social capital for the targeted rating level, there is a complementary method that involves decreasing the expected social capital gain of alternative rating levels by adding "unhelpful," which again results in a relative increase of expected social capital gain for targeted rating levels. To implement and quantify this approach, we add 10 more "unhelpfuls" for each 1- and 2-star rating to decrease raters' expected social capital gain for rating lower. Third, we note that the comparative additional social capital gain for rating at the targeted level is more prominent if we combine these methods. Therefore, we add 10 more "helpfuls" for

4- and 5-star ratings and 10 more "unhelpfuls" for 1- and 2-star ratings. Finally, we note the non-ordinal property of social capital gain, which suggests that adjusting expected social capital gain can be used to encourage raters to rate. Therefore, instead of increasing the rating valence, we assume the opposite goal of decreasing the rating valence by adopting the opposite approach of increasing the "unhelpfuls" for 4- and 5-star ratings by adding "helpfuls" for 1- and 2-star ratings. Additional approaches can be developed by allowing adjustments to the number of expected replies and decreasing expected RIRs. Due to space limitations, we do not extend the experiment to more diversified settings, although it would not be difficult to show the similarity and mirroring of effects compared to the experiment set forth here. Figure A.1 shows the output of those four experiments.

Figure A.1. Ratings Distribution with Adjusted Expected RIR

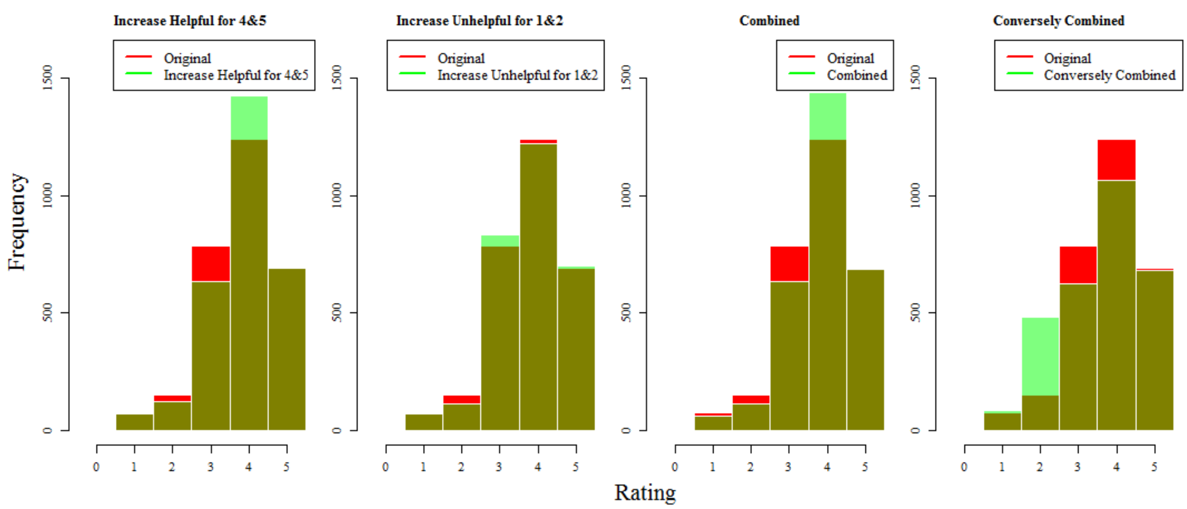

By increasing expected social capital for rating at a higher level, which is our targeted level, raters are encouraged to give more 4-star ratings as substitutes for 3- and 2-star ratings. The average rating is 3.863, which is significantly higher than the 3.791 in the original observations, suggesting the effectiveness of increasing ratings at the targeted level by increasing the corresponding expected social capital gain. Similarly, decreasing the expected social capital gain for alternative rating levels lowers rating volume for 1- and 2-star ratings, which are redistributed to 3- and 5-star ratings. Surprisingly, the volume of 4-star ratings decreases slightly in our experiment. However, the average rating valence is 3.804, which is higher than the original average rating valence. By combining the first and second experimental designs, we note that a redistribution of rating volume from low to higher levels becomes more obvious, especially the decrease of 1- and 2-star ratings, which leads to a stronger impact for rating valence, shown as an average of 3.877. By adopting a conversely combined approach, and, thus, a targeting level of a lower rating valence, we also change the redistribution, with a significant increment of 1- and 2-star ratings and a significant decrement of higher levels, such as 3-, 4- and 5-star ratings. Consequently, the average rating valence is decreased to 3.604, indicating the effectiveness of lowering the rating valence by adjusting the expected gain in social capital.

Table A.2. Transition Matrix When Expected RIR Manipulated

	1-star	2-star	3-star	4-star	5-star	1-star	2-star	3-star	4-star	5-star
	Increase Helpful for 4&5					Increase Unhelpful for 1&2				
1-star	0.08	0.12	0.26	0.42	0.12	0.47	0.01	0.26	0.22	0.04
2-star	0.07	0.10	0.34	0.41	0.08	0.04	0.49	0.28	0.14	0.05
3-star	0.03	0.06	0.27	0.49	0.14	0.02	0.03	0.71	0.18	0.07
4-star	0.02	0.03	0.21	0.52	0.22	0.01	0.01	0.13	0.76	0.09
5-star	0.01	0.02	0.12	0.44	0.41	0.00	0.01	0.08	0.16	0.75
	Combined					Conversely Combined				
1-star	0.04	0.14	0.27	0.42	0.14	0.07	0.28	0.31	0.22	0.12
2-star	0.04	0.11	0.35	0.40	0.10	0.11	0.27	0.23	0.29	0.10
3-star	0.03	0.05	0.26	0.52	0.13	0.04	0.23	0.26	0.32	0.14
4-star	0.02	0.03	0.21	0.52	0.22	0.02	0.17	0.20	0.40	0.22
5-star	0.01	0.02	0.14	0.43	0.41	0.01	0.04	0.16	0.38	0.40

For the offsetting effect among ratings redistribution, we additionally calculate the transition matrix for our experiment to avoid potential underestimates of the impact of the policy, as shown in Table A2. The upper-left transition matrix shows that more than 50% of the ratings were redistributed due to the additional expected positive RIRs. Most redistribution applies to neighboring rating levels; e.g., 3- and 5-star ratings neighbor 4-star ratings, and 4-star ratings neighbor 5-star ratings. Although migration to both lower and higher rating levels coexists in this transition, migration to a higher rating level generally exhibits a dominating role as compared to migration to a lower rating level, which is shown as the generally larger scales of inputs in the upper-right-hand side of the diagonal in the transition matrix as compared to the scales in the bottom-left-hand side. This is consistent with the histogram shown in Figure A.1, which displays the overall effect after oppositely directed transitions offset each other. The upper-right transition matrix, which represents redistribution, given additional expected negative RIRs for 1- and 2-star ratings, shows that the unchanged ratings account for 47% to 76% of the total distribution. For the remainder of the raters who redistribute, the transition exhibits a pattern similar to that of the upper-left-hand side, albeit with relatively less strength. This is consistent with our finding that "unhelpful" exhibits relatively weaker effects than does "helpful." The bottom-left-hand side, representing the redistribution when adding "helpful" for high ratings and "unhelpful" for low ratings are both applied, showing the relatively lower magnitude of diagonal elements and the higher magnitude of the non-diagonal elements, with a pattern similar to that of the upper-left and upper-right matrix. This suggests the stronger impact on ratings redistribution when the two methods are combined. In other words, adding "unhelpful" to lower ratings has a marginal effect in the same direction as that of the effect of adding "helpful" to higher ratings. Finally, when we combine both of methods conversely, we see the opposite effect: The scales of the elements on the bottom-left-hand side of the diagonals exhibit a larger magnitude than do those on the upper-right side, suggesting that more ratings were redistributed to lower levels. The magnitudes of the off-diagonal elements are as large as those of the bottom-left-hand matrix, and the overall proportion of ratings on the diagonal is similar to those in the bottom-left-hand matrix, suggesting a similar magnitude of effect as compared with the third experiment.

VITA

Jinyang Zheng is an assistant professor at the Krannert School of Management at Purdue University. His current research interests include sharing economy, online two-sided markets, online transportation networks, and machine learning in business analytics. His recent study found that online transportation networks tools can significantly improve the operational flexibility and efficiency of public transportation system; early-stage large-scale sales promotion can stimulate consumer learning about online transportation networks; the "conform or to be cast out" policy can significantly enhance the consumer welfare; WeChat contributes to the usage of the smartphone ecosystem without squeezing out the usage of the other apps. Prior to his Ph.D. study the University of Washington, he received the bachelor of science degree in statistics from Fudan University in Shanghai, China.