

ABSTRACT

Title of Dissertation: PRICING AND EMPLOYMENT ISSUES IN
THE PROVISION OF RIDE SERVICES

Ziwei Cao, Doctor of Philosophy, 2022

Dissertation directed by: Professor Michael O. Ball
Department of Operations Management

Professor P.K. Kannan
Department of Marketing

Recent years have witnessed the emergence and dramatic growth of platform businesses. This dissertation addresses two challenges of significance to ride service companies: 1) it investigates promotion effects on two-sided platforms; and 2) it models platform pricing and staffing strategies under a hybrid employment mode.

In the first chapter, I broadly discuss the new challenges faced by the ride service platforms in recent years and provide perspectives on related research questions.

In the second chapter, I study the network effects of different promotion methods in two-sided markets. Using data from a transportation-service platform, I specify a structural model that quantifies the respective promotional effects for price discounts and service upgrades. The results show that the primary effect of price discounts is to increase demand within the same service tier, whereas upgrades have stronger stickiness effects and spillover effects. Based on the estimates, I calculate the return on investment (ROI) and find that the ROI for upgrades is higher than that for

discounts. Our counterfactual analyses show that as the platform matures, the importance of upgrades increases while the importance of price discounts decreases. These results provide important managerial implications for platforms on how to design optimal promotions to grow their business.

In the third chapter, I model an on-demand platform that adopts a hybrid employment mode. This work is motivated by the recent public debate over the status of drivers for the major ride-hailing platforms as contract workers. My hybrid employment environment includes both contractors and full-time employees who receive a benefits package. In the hybrid model with driver control, drivers have the flexibility to decide how long to work and consequently whether to be an employee or a Contractor. Those who work over a certain number of hours will be classified as employees and receive a benefits package. The platform is a profit-maximizer and decides the optimal price based on the required benefit amount. As the benefit amount increases, the platform's profit decreases, which is consistent with strong gig company opposition to providing benefits. Moreover, I show that higher benefits make consumers and full-time drivers better off but decrease part-time drivers' welfare as well as overall social welfare. I consider a second model to better balance the platform's profitability and drivers' welfare. Under the hybrid model with platform control, the platform hires a limited number of full-time employees while guaranteeing that a minimum proportion of all work will be fulfilled by those employees. In this way, the profit loss due to the required benefits is capped, making this alternative policy a potentially viable solution from the platform's perspective.

PRICING AND EMPLOYMENT ISSUES IN THE PROVISION OF RIDE
SERVICES

by

Ziwei Cao

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2022

Advisory Committee:

Professor Michael O. Ball, Chair
Professor P.K. Kannan, Co-Chair
Professor Tunay Tunca
Assistant Professor Lingling Zhang
Assistant Professor Chenyu Yang
Professor David J. Lovell

© Copyright by
Ziwei Cao
2022

Table of Contents

Chapter 1: Introduction	1
Chapter 2: Lower Price or Better Service	6
2.1 Introduction	6
2.2 Literature Review	11
2.3 Data and Reduced-form Evidence	15
2.3.1 Empirical and Institutional Settings	15
2.3.2 Descriptive Statistics	18
2.3.3 Reduced-form Evidence	24
2.4 Model and Estimation	28
2.4.1 Consumer Decisions	29
2.4.2 Driver’s Decisions	32
2.4.3 Platform Decisions	36
2.5 Results	40
2.5.1 Parameter Estimates	40
2.5.2 Promotion Effectiveness	43
2.6 Counterfactual Analyses	45
2.6.1 Providing Drivers Bonuses	45
2.6.2 Marketing mix at Different Stages	47
2.7 Conclusion	49
Chapter 3: Hybrid Employment Modes for Ride-hailing Systems	52
3.1 Introduction	52
3.2 Literature Review	55
3.3 Model	58
3.3.1 Contractor Only: Baseline	60
3.3.2 Blended Workforce: DC Model	61
3.3.3 Blended Workforce: PC Model	64
3.4 Numerical Study	67
3.4.1 Baseline vs. DC Model	67
3.4.2 DC vs. PC Model	69
3.5 Conclusion and Future Research	74
Appendices	78
Bibliography	88

Chapter 1: Introduction

Platforms are prevalent in today's economy and have transformed many traditional industries, such as Uber in the transportation service industry and Airbnb in the room/home rental industry. Unlike conventional services, platforms through pricing and other decisions, must actively manage both: the demand and supply sides. For example, ride-hailing companies such as Uber and Lyft match riders and drivers and must balance rider demand and driver supply through their pricing policies. Retailing platforms like Amazon and eBay connect buyers and sellers. Food delivery services like Grubhub and UberEATS link restaurants and eaters. In this dissertation, I will focus on ride service platforms.

Since the demand and supply are dynamically interrelated, attracting consumers and retaining drivers are both critical for the platform's profitability. To attract users, platforms typically provide incentives through marketing promotions. Price discounts and free upgrades are among the platforms' most common promotional practices. Discounts directly reduce the amount charged to consumers, while product and service upgrades—often in the form of trials—enable customers to try higher-tier products or services for the same price. The effects of discounts and upgrades have been studied extensively in marketing research. Discounts, in general, have a positive effect on sales, especially for price-sensitive customers. Upgrades can also increase sales by increasing awareness and reducing quality uncertainty. In two-sided markets, high demand attracts supply, and high supply draws demand, forming the so-called “network effect.” Promotions not only affect demand by attracting more consumers, but also have a spillover effect on the supply side. Moreover, this spillover effect on drivers will lead

to further changes in demand. In this way, understanding the influence of various promotions by considering the network effects is prominent for platforms. Extant literature has examined the promotion effects largely on the demand side, without explicitly considering the spillover effect on the supply side. Few empirical studies have explicitly examined how the network effect affects promotions. Thus, previous findings on demand-side promotion effects are insufficient to guide decision-making for platforms. The second chapter aims to fill this important gap. In this research, we ask three research questions: (1) With network externalities, how do discounts and upgrades affect demand and supply differently? (2) How does the return on investment (ROI) for these two promotions differ? And (3) What factors determine the marketing mix of promotions?

On the supply side, platforms can consider different employment modes to incentivize drivers. For example, based on both survey and administrative data, Hall and Krueger (2018) shows that drivers who partner with Uber appear to be attracted to the platform largely because of the flexibility it offers, the level of compensation, and the fact that earnings per hour do not vary much with the number of hours worked. On the other hand, a full-time driver-employee would receive protections and employment benefits such as unemployment insurance and health insurance. An employer usually has direct control of its employees and makes centralized staffing decisions. In the meantime, the guaranteed protections help increase drivers' loyalty and stickiness to the platform. However, part-time contractor-drivers would work less while valuing flexibility and freedom: they like to choose when and how long to work instead of sticking to fixed schedules. Traditional ride services adopt the employee-only mode.

For example, taxi drivers are hired as employees and receive a benefits package. Their shifts are predetermined by the platform. The new emerging ride-hailing companies like Uber and Lyft classify their drivers as independent contractors. Drivers decide when and where to open the app but receive no protection from the platform. A debate emerged as to whether to classify independent-contract drivers as employees in recent years. Policymakers seek labor protections for drivers while platforms avoid sacrificing their profit margin. Therefore, how to balance flexibility and benefits for drivers with the profitability for platforms is a fundamental question for the future of the ride service industry. Prior research has focused on cases where drivers self-schedule, and on-demand platforms only hire independent contractors. However, facing new challenges, previous business models may not be applicable. Policies such as AB5 and Proposition 22 feed a growing literature that explores hybrid employment modes: platforms hire both employees and independent contractors. In the third chapter, I model an on-demand platform that adopts a hybrid employment mode that allows both contractors and full-time employees. The goal is to devise a policy:

- 1) that accommodates drivers who value flexibility and wish to remain contractors and drivers who want the benefits of full-time employment and addresses the related ongoing public debate.
- 2) that enables the ride-hailing platforms to maintain reasonable profit margins and viable business model.

Transportation service providers usually employ one of two modes of operation: on-demand service and reservation-based service. For example, I evaluate the promotional effects using data from a reservation service platform in the second chapter

and focus on the hybrid employment mode for on-demand platforms in the third chapter. Convenience and flexibility may be the most important characteristics contributing to the popularity of on-demand services. Moreover, with the new sharing economy business model, price-sensitive customers can view pricing for alternative service providers in real-time and choose appropriately. Unlike on-demand services, reservation services have their own advantages: once a reservation is made, the service will be guaranteed, so a customer has less uncertainty about the fulfillment of her future demand compared to an on-demand service. Moreover, a customer does not need to incur extra waiting time for identifying and arranging for the service. Reservations also benefit platforms: companies gain more information about the future demand and have a longer time to coordinate demand and supply.

While both systems are prevalent nowadays, it is worthwhile to discuss the revenue management problem under competition between reservation-based and on-demand services. Potentially, several research questions could be addressed: On the demand side, how do two systems make pricing decisions? In current practice, reservation systems usually adopt a static pricing strategy: prices are predetermined and uniform for consumers. However, on-demand platforms adopt dynamic pricing, and the prices may vary accordingly. Meanwhile, consumers have different valuations for different service types: Those who tend to plan and schedule in advance value the reservation system more than the on-demand service. Also, with the reserved orders, usually drivers will wait for customers before their departures and customers can thus avoid the waiting cost. However, customers who have higher uncertainty about their future trips may prefer the on-demand service to avoid the penalty of booking in advance but

end up with a no-show. Of course, it is possible that some customers are indifferent between the two systems. Then how to compete in terms of prices and service qualities facing different types of customers is critical for platforms. At the same time, retaining drivers also plays an important role: On the supply side, platforms need to decide how to pay their drivers and the optimal capacity level. For example, Uber drivers are paid by job currently. When delays caused by consumers happen, drivers need to incur the waiting cost by themselves, which may lead to reduced hourly earnings. However, if platforms pay drivers by the hour, drivers are guaranteed a certain hourly rate but how to incentivize drivers becomes challenging. Meanwhile, a hybrid employment mode can be adopted by two different services: the reservation-based platform with less demand uncertainty may consider hiring more full-time drivers, and the on-demand platform may employ more part-time drivers due to the fluctuating demand.

Chapter 2: Lower Price or Better Experience

2.1 Introduction

Platforms are prevalent in today's economy and have transformed many traditional industries such as Uber in the transportation service industry and Airbnb in the home-sharing industry. A defining characteristic of a platform is that the demand and supply are dynamically interrelated: high demand attracts supply, and high supply draws demand, forming the so-called "network effect." However, the network effect can also be a double-edged sword creating the well-known "chicken-and-egg" conundrum: a platform needs demand to compete for supply and needs supply to compete for demand. The question becomes which side should be increased first. This is a common challenge for all platforms, especially for new entrants and underdogs. Thus, understanding the effectiveness of various user-growth strategies takes on special importance for platform research.

To attract users, platforms typically provide incentives through marketing promotions. Price discounts and free upgrades are among the most common promotional practices on platforms. Discounts directly reduce the amount charged to consumers while product and service upgrades—often in the form of trials—enable customers to try higher-tier products or services for the same price¹. The effects of discounts and upgrades have been studied extensively in marketing research. Discounts, in general, have a positive effect on sales, especially for price-sensitive

¹ In subsequent analysis, we use "products" and "services" interchangeably.

customers. Upgrades can also increase sales by increasing awareness and reducing quality uncertainty.

Extant marketing literature has examined the promotion effects largely on the demand side, without explicitly considering the spillover effect on the supply side. For many platform businesses, a shock to demand would cause a chain reaction to supply, which would lead to further changes in demand. Although this reinforcing relationship between the demand and supply (i.e., the network effect) is prominent for platforms, few empirical studies have explicitly examined how the network effect affects promotions. Thus, previous findings on the demand-side promotion effects are insufficient to guide decision-making for platforms. The current study aims to fill this important gap. In this research, we ask three research questions: (1) With network externality, how do discounts and upgrades affect demand and supply differently? (2) What is the return on investment (ROI) of these two promotions, respectively? And (3) What factors determine the marketing mix of promotions?

To answer these questions, we work with a transportation service platform that is a leading player in a major Asian market. The platform matches riders (i.e., “customers”) and individual drivers (i.e., “suppliers”) to provide transportation services. The service is similar to Uber and Lyft, except that rides are specialized to pick-up and drop-off riders at airports and train stations. Our data cover all transactions from the platform’s largest markets over 20 months, with more than a million unique riders, 60,000 drivers, and approximately 74,000 reservations per month. On the demand side, the platform provides three service tiers: basic, premium, and luxury. The higher the tier, the more upscale the vehicle. The platform frequently offers two promotions to customers: price

discounts and upgrades. With the former, a customer can enjoy a ride at a lower cost; for the latter, the customer is bumped to a higher service tier at no additional charge. On the supply side, the platform works with two types of drivers: in-network and out-of-network. In-network drivers have a contract with the platform and agree to serve any orders dispatched to them. In turn, the platform dispatches call to these drivers first. The out-of-network drivers do not have a contract and only serve when in-network drivers are unavailable. In terms of promotions, the coupons (i.e., price cuts) can be applied to rides from either type of driver, but the platform only uses in-network drivers for upgraded rides.

We develop a structural model to examine the effects of discounts and upgrades, including three models to capture the customer's decision, the driver's decision, and the platform's decision, respectively. On the demand side, we specify a two-stage model: (1) a customer first decides whether or not to make a reservation, and (2) she chooses which tier to book. Because the customer can use a coupon for any tier, price cuts are entered in the first-stage reservation model. At the second stage, the customer's tier choice depends on the tier-specific price, preference, and whether or not the customer has had previous experience with this tier. The tier-specific number of drivers is also included in the demand model to explicitly capture the network effect. On the supply side, we model the driver's retention decision as a function of the driver's expected wage on the platform, which captures the cross-side network effect from demand to supply. In the platform model, we assume that the platform maximizes its profits by strategically determining the promotion levels.

Our parameter estimates confirm that cross-side network effects are positive and significant on this platform, verifying that it is important to include the spillover effect in the model. Based on the parameter estimates, we quantify the profit gains of providing discounts and upgrades and calculate the respective ROI. We find that the ROI for upgrades is higher than that for discounts, perhaps because upgrades not only reduce the customer's uncertainty about the higher-margin services but they also directly help driver retention by increasing their wages. Interestingly, the ROI for upgrading from a basic to a premium service tier is higher than the ROI for upgrading from a premium to a luxury service. This is jointly determined by the fact that a free trial can lead to high incremental demand for the premium tier and that it is relatively costly to replace a premium-tier driver when attrition occurs.

How can the platform improve its user-growth strategy based on our findings? To answer this question, we conduct two counterfactual analyses. In the first analysis, we allow the platform to consider distributing incentives on both sides of the platform rather than just focusing on the demand side as has been done in the past. Our results show that a bonus for drivers has a positive ROI and can improve platform profits. This indicates that giving direct monetary incentives to drivers, rather than indirectly influencing drivers through orders, is a viable strategy to consider. In the second analysis, we examine how the platform should adjust the marketing mix of discounts and upgrades at different stages of its lifecycle. We find that at the early stage, the effectiveness of discounts is close to that of upgrades. As more customers experience the basic and premium service tiers, the value of providing discounts diminishes. Thus,

discounts are more effective in the early stage, and upgrades become more effective as the platform gains greater market share.

Our paper provides two primary contributions to the literature. First, we extend the promotion literature to platform settings. Our empirical model provides a framework for examining the effectiveness of price discounts and upgrades in a two-sided market. Our analysis indicates that discounts and upgrades affect the demand and supply in different ways. Compared with discounts, free upgrades not only affect consumer demand but also have a more direct impact on the supply side. Discounts increase the demand, which positively affects supply via the cross-side network effects. Upgrades shift demand to higher tiers, increasing the wages of higher tier drivers and the likelihood of these drivers to remain on the platform. Our method can also be applied to quantify the ROI of marketing promotions for companies with network effects. Second, our study contributes to research on platforms' user growth strategies. We demonstrate that it is important to unbox the mechanisms through which each promotion method grows the customer base and supply base. In our setting, upgrades directly affect both the demand and supply: it increases demand by reducing quality uncertainty and attracts drivers by increasing wages. Thus, the ROI is higher than that for discounts in the early stage of the platform. Our paper is the first empirical work that examines multiple alternate strategies for two-sided platforms. Our research provides insights on how a platform should adjust its optimal marketing mix at different stages of its lifecycle, which is novel to the literature and generates actionable managerial implications in practice.

2.2 Literature Review

Our research is related to two streams of literature: promotion effects and two-sided markets. In this section, we briefly discuss relevant studies and present the contributions of our study.

Promotion Effects

Sales promotions have been studied extensively in marketing research for decades. Promotions in the form of price discounts are found to affect consumers' purchase behavior directly. In their comprehensive review paper, Ailawadi and Gupta (2014) summarizes findings on various promotions, including coupons, buy-one-get-one-free, free samples, and rebates, and categorize consumers' responses into immediate and long-term responses. Our paper studies two of the most common promotion methods: coupons and free samples (free upgrades in our context). Coupons can increase sales through an advertising exposure effect or redemption (Bawa et al. 1997; Neslin 1990; Reibstein and Traver 1982; Ward and Davis 1978). Neslin (1990) estimates the incremental sales per coupon redemption and finds that the coupons disproportionately attract loyal customers. Bawa et al. (1997) also considers the joint effects of coupon attractiveness and coupon proneness on redemption. Other studies have examined the impact of coupon characteristics on coupon redemption rates. For example, a higher coupon face value is found to lead to a higher redemption rate (Reibstein and Traver 1982; Ward and Davis 1978).

Free samples are another common promotional method used by many businesses. Heiman et al. (2001) decomposes the sampling effort into two parts: the immediate effect on sales and the long-term effect on goodwill (i.e., a learning effect). Bawa and

Shoemaker (2004) also empirically shows that free samples could produce a more lasting effect on sales than coupons. Using analytical modeling and field experiments, Li et al. (2019) studies whether to provide free samples and determine the optimal design of the samples. In a similar service upgrade setting to the current study, Sudhir and Yang (2018) examines free samples (i.e., upgrades) in car rental services and find that upgrades induce greater loyalty and thus have a spillover effect on demand across tiers with different profit margins. However, our study is different from Sudhir and Yang (2018) in two aspects: (1) we jointly model the impacts of upgrades on both demand and supply, and (2) we aim to identify the different mechanisms through which price discounts and upgrades affect platform profits, respectively.

Two-Sided Markets

Early work provides a theoretical framework on competition and pricing strategies in the presence of cross-side network effects (Caillaud and Jullien 2003, Rochet and Tirole 2003; Armstrong 2006). With the increasing importance of online platforms, empirical work related to two-sided markets has been growing rapidly. Empirical research has estimated the network effect and examined the impact on pricing in different industries. Clements and Ohashi (1994) and Dubé et al. (2010) shed insights on pricing structures in video gaming. Chiara et al. (2020) shows how digital platforms balance network effects and differentiation. Chu and Manchanda (2016) quantifies the asymmetrical cross-side network effects in online retailing. Ming et al. (2019) studies ride-hailing services and it is the first empirical work that explicitly models both consumer and driver decisions. They explore the effects of surging pricing and other regulatory policies on both ride-hailing market and the traditional taxi market. Built

upon Ming et al. (2019), our paper addresses a related but different question: how and to what extent different promotions affect the demand and supply in a two-sided market, respectively. In addition to modeling the decision-making process for consumers and drivers, our paper also explicitly takes into account the platform's strategic decision.

The existence of network effects on platforms is critically important for platform growth. When a platform enters the competition, its smaller consumer base makes it unattractive to suppliers. The smaller supplier base is also unappealing to consumers. This is famously known as the “chicken-and-egg” conundrum (Parker et al. 2016). Given this challenge, platforms are incentivized to use promotions to grow their user bases and leverage the positive feedback loop between demand and supply. Promotions such as price discounts and free upgrades could potentially have a greater impact on platform businesses than on traditional businesses. However, empirical research on this topic is limited with just a few recent exceptions. Zhang et al. (2020) examines how some sellers' decisions to offer a specific type of price promotion affected consumers' interaction with other sellers on the same platform. They document that promotions can spillover to sellers without promotions. Kabra et al. (2017) compares the effect of incentives given to passengers and drivers in a ride-hailing marketplace, respectively, and find that it is more efficient to give incentives to passengers in the short run and to drivers in the long run.

There is one significant limitation remaining: much of the existing literature has only studied price discounts. None of these papers have explicitly examined the mechanisms of more than one promotion, although platforms commonly use multiple

forms of incentives to grow their user bases. Marketing literature has established that free upgrades allow consumers to experience an unknown service, which could lead to higher user loyalty and stickiness (Bawa and Shoemaker 2004). The network effect could further magnify such an effect in a two-sided setting. Thus, platforms need to understand exactly how each promotion method works and when to use which. The few studies examining other promotion methods either have not yet sufficiently addressed the network effects or have only studied one promotion. Albuquerque et al. (2012) measures the value of promotional activities, including price discounts on an online platform where users can create and purchase content. In a context similar to ours, Zheng et al. (2020) models how sales promotions affect drivers' willingness to use the app but only incorporate drivers' forward-looking behavior, and not cross-side network effects. Sudhir and Yang (2018) models an experience effect due to free samples but also does not include a network effect. Our study investigates two of the most common promotions on two-sided markets by modeling the consumers', drivers', and the platform's decision-making processes. On the demand side, we include stage dependency to capture the experience effect due to upgrades. We also model the cross-network effects on both the demand and supply sides. Table 2.1 summarizes our contributions compared to extant work.

Table 2.1 Summary of Relevant Studies on Two-sided Markets

Research	Context	Research Focus	Promotion Strategy	Quantify the ROI of promotions	Model Platform's Decision
Chu and Manchanda (2016)	Retailing	Quantify network effects	No	No	No
Kabra, Belavina and Girotra (2017)	Ride-hailing	Effectiveness of incentives	Discounts	No	No
Sudhir and Yang (2018)	Car Rental	The value of free upgrades due to state-dependence	Free upgrades	No	Yes
Our Research	TNC	Effectiveness of promotions; when to provide which	Discounts and free upgrades	Yes	Yes

Notes: TNC=transportation network company

2.3 Data and Reduced-form Evidence

2.3.1 Empirical and Institutional Settings

We work with one of the largest transportation-service platforms in Asia. Customers use the platform to book a ride for airport and train station pickups or drop-offs. Reservations can be made via the app or on the platform's website, with the former being the primary booking channel. During the booking process, after a consumer enters the information such as flight number, time, and destination, different booking options with respective price would be shown accordingly. The available options on the platform include multiple tiers of services—economy, comfort, business, luxury—and the specific car types are also provided. For tractability, we group those service types into three main service tiers: basic, premium, and luxury, according to different car models and therefore prices.² For each reservation, the platform dispatches it to

² During our data collection periods, drivers remained within the same tier and never switched across tiers.

drivers from the chosen tier. When the order is accepted, the customer receives a notification with the driver's phone number, car model, and plate number. On our focal platform, consumers usually make their reservations ahead of the trip. Based on the data, the duration between the booking time and the service time ranges from one day to 8 months, with over 90% of the consumers making the reservation within one week of their trip. This may be because consumers prefer less uncertainty and would not want to make a reservation too early ahead of time. Customer can of course cancel the reservation at any time without penalty, as long as it is made at least two hours prior to the departure time.

On the supply side, there are two types of drivers on the platform: in-network and out-of-network. The in-network drivers have a contract with the platform and agree to work exclusively with the focal platform while the out-of-network drivers do not have such an exclusivity contract. In particular, in-network drivers commit a fixed shift so that the platform can assign them orders directly. The platform therefore treats these two types of drivers with different accommodation priorities. When receiving a reservation, the platform first dispatches the order to in-network drivers. Only when in-network drivers are unavailable does the platform broadcast the order to out-of-network drivers. In-network drivers renew their contracts at the beginning of each month, while out-of-network drivers are free to opt in or opt out of the system. Since the platform is actively hiring, there is no restriction on the number of in-network drivers. Out-of-network drivers are also allowed to start a contract and work exclusively for the platform. The platform has a direct control for in-network drivers and an indirect control for out-of-network drivers: orders are assigned to in-network drivers and

usually the dispatched orders cannot be declined; the platform broadcasts the excess orders to out-of-network drivers who can decide whether to accept the order based on their availability.

Drivers' retention decision is based on their expected revenue, which is correlated to the expectation of future demand. To better help drivers make decisions, the platform usually shares the demand and promotion data in the past six months. With historical data, drivers can have a better estimation for the future demand.

Drivers are paid a fixed revenue share per ride. The commission rate for out-of-network drivers is higher than that for in-network drivers, and the rate is constant across service tiers. The driver's service tier is predetermined by the car model, which is fixed during the contract. Even though the platform offers promotions to consumers, drivers are paid according to their tier-level compensation rate. For example, a premium tier driver who takes an upgraded order from the basic tier receives the compensation based on the price of premium tiers. Note that a coupon can be applied to any service tiers, but upgrades apply only to in-network drivers.

The platform actively employs two promotional methods to grow its customer base: price discounts and service upgrades. The optimal level of promotions at the market level are determined by the platform at the beginning of each month. Price discounts are implemented as coupons, which allow customers to enjoy a direct price deduction during booking, e.g., 5 yuan off or 10% off the total cost. In contrast, service updates allow customers to enjoy a better service without additional costs. For example, a customer who makes a reservation for the basic service tier could be upgraded to the premium- or luxury-car service. When this occurs, the customer still pays the basic-tier

rate, which is lower than the price charged by the higher-tier driver, and the difference is absorbed by the platform as a marketing cost. Consumers know the coupon face value before they search and make the reservation. After the reservation is made, the platform starts to dispatch the order, and consumers will be informed of the driver's information once the order has been successfully dispatched. The upgrade is unknown to consumers until the dispatch is finalized.

Per our discussions with the management team, during our data collection period, the focal platform did not employ any personalization in their promotion strategies. In other words, the coupon distribution (in terms of frequency and discount depth) and the upgrade opportunity are an endogenous decision of the platform based on their profit maximization objective but not based on customer characteristics or transaction history.

If the customer uses a coupon or receives an upgrade, the platform receives the revenue based on the reservation rate minus the coupon (if applicable) but pays the driver according to his or her tier-level compensation rate. When the revenue is lower than the platform's payment to the driver, the profit for the platform is negative, which can be thought of as the platforms' marketing cost. Note that a coupon can be applied to any service tiers, but upgrades are applicable only to in-network drivers.

2.3.2 Descriptive Statistics

Customer reservations and tier choice

In this section, we describe our data and present descriptive statistics. Our data include all the reservations from the platform's thirteen largest markets over a period of 21 months (from January 2017 to September 2018). Each reservation includes the

information on four aspects: the order fulfillment, the rider, the driver, and the payment. The order fulfillment information includes the pick-up and drop-off locations, ride time, fulfilled service tier, and, if the order is canceled, the reason(s) for cancellation. The rider information includes the flight or train information, requested service type (which would be different from the fulfilled service tier in the case of an upgrade), and reserved time. The driver information includes the driver's ID, car type, in-network status, and the time when the driver accepted the order. The payment information includes the estimated price (which the customer sees when making the reservation), the final ticket price for the ride (before a discount), the actual payment (after a discount), and the discount depth (if applicable).

Table 2.2 presents the summary statistics for reservations by service tier. Approximately 52% of the reservations were made for the basic tier, 38% for the premium tier, and the remaining 10% for the luxury tier. The average ticket price is ¥120.0 (SD = 25.6), ¥136.0 (SD = 24.6), and ¥191.3 (SD = 43.0) for the basic, premium, and luxury tiers, respectively.³ Reservations for the basic or premium tier have a positive probability of being upgraded to a higher tier. On average, 18.33% of the basic-tier reservations were promoted to premium, and 7.96% of the premium-tier reservations were upgraded to luxury. In theory, a basic-tier reservation could be upgraded to the luxury tier, but such upgrades are negligible in our data.

In terms of coupons, coupon incidence is the percent of reservations that used a coupon. In our data, more than 60% of the orders applied a coupon, regardless of the service tier. The average coupon depth ranged from 17.40% off (the ticket price) for

³ The means are equivalent to USD \$18.50, \$21.00, and \$29.60 for the three tiers, respectively.

the luxury tier to 25.65% off for the basic tier. The average face value of coupons is similar across tiers, the mean being approximately ¥32.0.

Table 2.2 Descriptive Statistics of Reservations by Service Tier

	Service Tier		
	Basic	Premium	Luxury
# of Bookings	2,962.1 (1,913.4)	2,420.8 (2263.7)	604.2 (484.7)
Fraction of Bookings (%)	52.2 (14.1)	37.9 (13.9)	9.9 (3.8)
Ticket Price (¥)	120.0 (25.6)	136.0 (24.6)	191.3 (42.0)
Upgrade Rate (%)	18.33 (14.22)	7.96 (9.04)	0 (0)
Coupon Incidence Rate (%)	65.20 (11.61)	68.21 (9.69)	66.26 (11.55)
Coupon Depth (%)	25.65 (7.79)	26.04 (6.67)	17.40 (7.00)
Coupon Face Value (¥)	32.11 (10.11)	32.35 (10.97)	32.23 (14.13)

Note: The unit of observation is a market-month combination. The numbers are mean and standard deviation (in parenthesis).

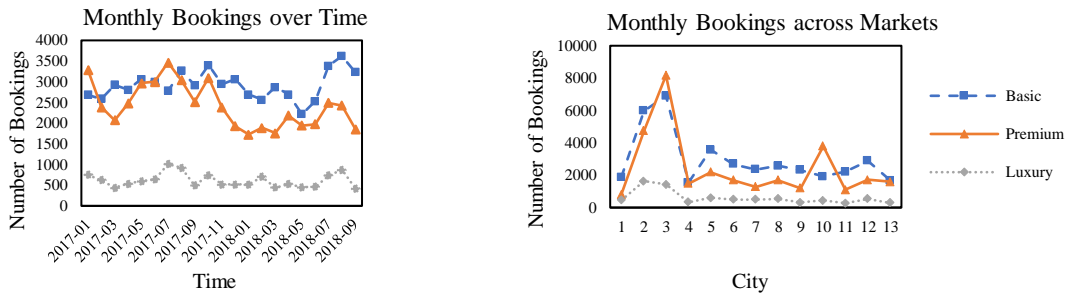


Figure 2.1 Reservations by Month and Market

In Figure 2.1, we depict the number of reservations for each tier over time (the left panel) and across markets (the right panel). We find that there is substantial seasonal variation over time and geographical variation across markets, indicating that it is critical for our model to account for time and market heterogeneity. Furthermore, we plot the average coupon value by tier in Figure 2.2 and the update rate in Figure 2.3,

which also confirm the data variation by tier, month, and market. The within-market panel data variation per tier plays an important role in identifying our parameters of interest, which will be explained in more detail in the Model and Estimation Section.

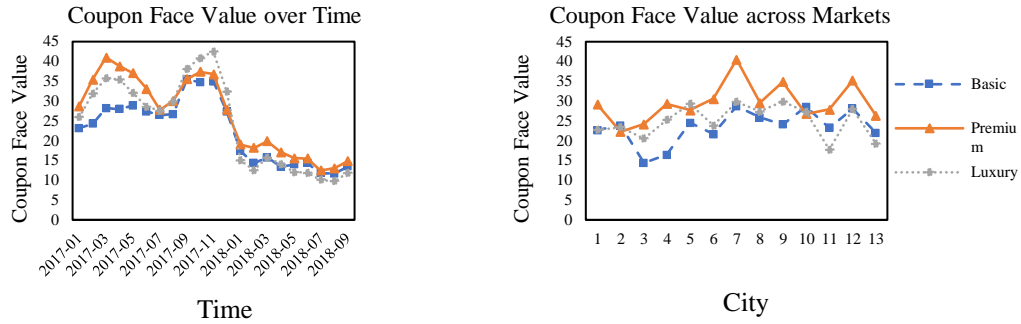


Figure 2.2 Average Coupon Face Value by Tier




Figure 2.3 Upgrade Rate by Tier

Lastly, we compare the marketing cost of providing upgrades and coupons from the platform’s perspective. An upgrade from basic to premium tier is equivalent to a coupon with a 16.0% price discount, and an upgrade from premium to luxury tier is equivalent to 28.9% off. The average coupon depth for premium and luxury are 26.0% and 17.4%, respectively (see Table 2.2). Thus, from the marketing cost perspective, a free upgrade is less costly for basic-tier reservations but more costly for premium-tier reservations.

Driver behavior

In this section, we present descriptive statistics for the in- and out-of-network drivers. Table 2.3 presents the summary statistics for driver-related variables. As expected, there are more lower-tier drivers than higher-tier drivers on the focal platform. There are also more out-of-network drivers than in-network drivers: on average, each market has 80.54 (SD=31.19) in-network drivers and 1,051 (SD=934.25) out-of-network drivers who ever fulfilled an order in our data. The in-network drivers fulfilled more orders than out-of-network drivers did. On average, an in-network driver provides 30.51 rides (SD=8.85) per month, in contrast to the average of 2.24 rides (SD=1.18) for an out-of-network driver. This is consistent with the fact that the platform gives priority to its in-network drivers. Figure 2.4 depicts the composition of in- and out-of-network drivers from different tiers and the average rides they provide each month.

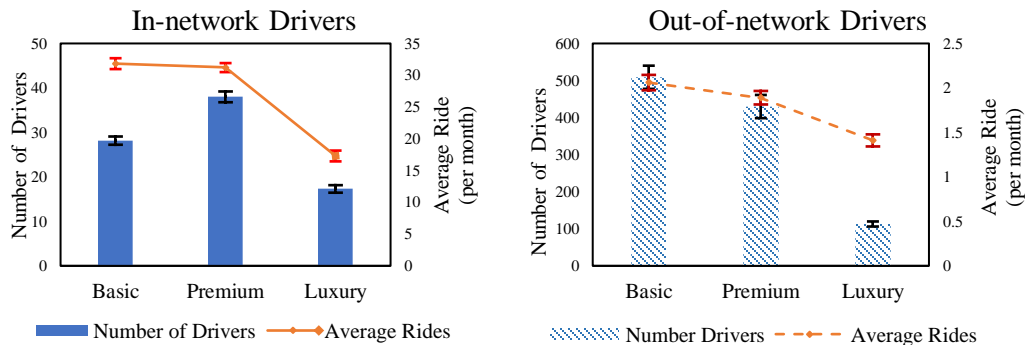


Figure 2.4 Number of Drivers and Average Rides per Driver by Tier⁴

⁴ Note: The height of the bar represents the average number of drivers per market per month for each tier, with the error bar representing 95% confidence interval. The line chart depicts the average number of rides and the 95% confidence interval for each tier.

Table 2.3 Descriptive Statistics of Driver Variables by Service Tier

	In-network Drivers				Out-of-network Drivers			
	Basic	Premium	Luxury	All	Basic	Premium	Luxury	All
# of Drivers	28.15 (13.50)	37.99 (17.74)	17.29 (12.34)	80.54 (31.19)	508.74 (457.21)	429.74 (456.81)	112.89 (99.63)	1,051 (934.25)
Average Rides (per driver per month)	31.82 (12.51)	31.20 (10.37)	24.66 (12.41)	30.51 (8.85)	2.06 (1.26)	1.89 (1.11)	1.41 (0.99)	2.24 (1.18)
Earnings (per ride)	84.01 (17.92)	95.20 (17.22)	133.91 (29.40)	122.91 (24.05)	96.00 (20.48)	108.80 (19.68)	153.04 (33.60)	140.47 (27.48)
Retention Rate (%)	65.82 (17.08)	67.89 (13.45)	68.67 (17.01)	67.40 (11.85)	42.86 (16.37)	44.19 (16.06)	45.04 (17.19)	43.66 (15.21)
# of Newly Added Drivers	8.63 (7.96)	12.30 (9.27)	5.82 (5.74)	26.74 (17.61)	227.30 (220.39)	255.42 (318.52)	65.07 (68.79)	547.80 (464.01)

The platform charges a fixed commission rate: 30% from the orders fulfilled by in-network drivers and 20% from those fulfilled by out-of-network drivers. The commission rate is constant across service tiers. It is straightforward to see that the average earning per ride for in-network drivers is 84.01 (SD=17.92) for the basic tier, 95.20 (SD=17.22) for premium, and 133.91 (SD=29.40) for luxury, which is proportional to the ticket price for each corresponding tier.

A key characteristic of ride-share platforms is that drivers can actively decide whether to stay on the platform. A stable driver base is critical for sustaining the network effect. Thus, we examine driver retention and turnover behaviors. On our platform, the in-network drivers are more likely to retain than the out-of-network drivers: the overall retention rate is 67.40% (SD=11.85%) for in-network drivers and 43.66% (SD=15.21%) for out-of-network drivers. Across the service tiers, drivers from

the premium and luxury tiers have a slightly higher retention rate than those from the basic tier.

2.3.2 Reduced-form Evidence

To motivate our structural model, it is important to understand on how the two promotion methods (i.e., price discounts and upgrades) are associated with consumers' choices and drivers' decisions. In this section, we present the ordinary least square (OLS) regression results and discuss some descriptive evidence.

Consumer decisions

For individual consumers, two decisions are of particular importance: (1) to return to the platform, and (2) to switch to a higher service tier. We measure a customer's returning decision using the number of repeated reservations for each customer. Among the over one million unique users on the platform, roughly 22% made multiple reservations during our data collection period, among whom the majority had less than five transactions (see Figure 2.5).

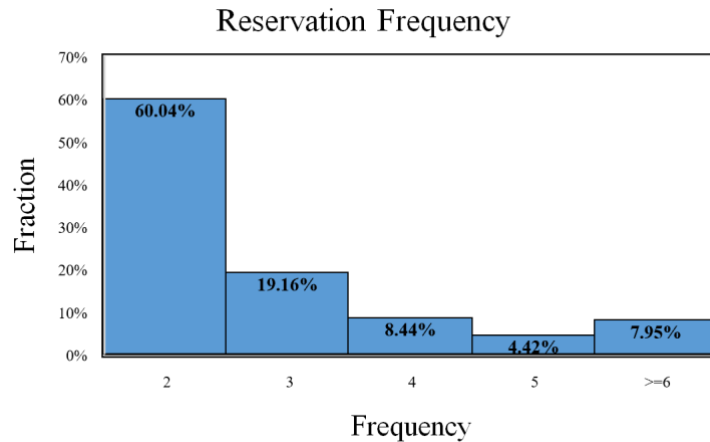


Figure 2.5 Reservation Frequency for Individual Customers

The switching behavior refers to whether the customer voluntarily switched to a higher tier in future transactions. For example, if customer i in market m booked the

basic tier before transaction f and the premium tier at transaction f , s/he is categorized as $SwitchHigh_{imf} = 1$. If s/he consistently booked the same tier, $SwitchHigh_{imf} = 0$. For customers without multiple reservations, $SwitchHigh_{imf}$ is undefined. Customers' transaction frequency and switching status are then regressed on the promotion variables using the following specification:

$$Frequency_{im} = \beta_0 + \beta_1 Upgrade_{im} + \beta_2 Coupon_{im} + \Phi_m + \epsilon_{im} \quad (2.1)$$

$$SwitchHigh_{imf} = \beta_0 + \beta_1 PastUpgrade_{im,f-1} + \beta_2 Coupon_{imf} + \beta_3 PastCoupon_{im,f-1} + Trip_{imf} + \Phi_m + \epsilon_{imf} \quad (2.2)$$

In Equation 2.1, $Upgrade_{im}$ is defined as 1 if customer i in market m ever received an upgrade and 0 otherwise; $Coupon_{im}$ is the average discount amount (¥) that i used across all bookings; and Φ_m is the market fixed effect. In Equation 2.2, the dependent variable is the switching status at transaction f of customer i . Thus, variable $PastUpgrade_{im,f-1}$ equals 1 if customer i upgraded at the previous transaction $f - 1$ and 0 otherwise; $Coupon_{imf}$ is the coupon amount that i applied for the current transaction f ; $PastCoupon_{im,f-1}$ is the coupon amount for the transaction $f - 1$; $Trip_{imf}$ captures the trip characteristics, such as the trip distance and whether it happens during weekdays; and Φ_m is again the market fixed effect.

The regression results are presented in Table 2.4. We find that having upgrades is positively associated with customers' transaction frequency and the switching status. Customers who upgraded in the past are not only associated with more reservations, but also are more likely to switch to a higher tier in the future. Not surprisingly, a higher price discount is positively associated with reservations. However, the influence of price discounts is mixed on the likelihood of switching: the past coupon has a positive

effect while the current coupon’s effect is negative, indicating that receiving a deeper discount motivates consumers to stay with the current tier while deeper discounts in the past is positively associated with the odds of customer switching to a better tier.

Table 2.4 OLS Regression Results on Customer Decisions

	Frequency Est (SE)	SwitchHigh Est (SE)
Upgrade/PastUpgrade	0.734*** (0.004)	0.823*** (0.021)
Coupon	0.117*** (0.001)	-0.003*** (0.000)
PastCoupon		0.014*** (0.000)
Distance		-0.010*** (0.001)
Weekday		-0.013 (0.010)
Constant		-2.953*** (0.049)
Market Fixed Effect	Yes	Yes
Observations	1,128,458	257,683

*** p<0.001; ** p<0.01; * p<0.05; + p<0.1

The results in Table 2.4 seem to suggest that price discounts can help retain customers on the platform, but they are not effective in promoting higher-tier services (which have higher margins and thus are more profitable for the platform). In contrast, service upgrades show a positive correlation with tier promotion as well as customer retention. These correlations provide suggestive evidence that these two promotion methods may have different managerial implications for the focal platform, which we will formally model in the Model and Estimation section.

Driver decisions

On the driver side, an important question is whether a driver is retained on the platform, because a stable supply of drivers is critical to attract demand and sustain the network effect for the platform. This analysis is only relevant for the in-network drivers

because out-of-network drivers serve the platform only when there is an excessive demand.

We use variable Retain_{jkmt} to denote whether driver j of tier k remained on the platform in market m at time t . We model the retention behavior as a function of several factors: (1) the driver individual fixed effect α_j ; (2) the wage that driver j earned in the past, $\text{PastWage}_{jkm,t-1}$; (3) the total demand at time t ; (4) the tier fixed effects T_k , and (5) the market-time fixed effects Φ_{mt} . The specification is as follows:

$$\text{Retain}_{jkmt} = \alpha_j + \xi_1 \log(\text{PastWage}_{jkm,t-1}) + \xi_2 \log(\text{Demand}_{kmt}) + T_k + \Phi_{mt} + \epsilon_{jkmt} \quad (2.3)$$

Estimates of Equation 2.3 are reported in Table 2.5. Results show that after controlling for the driver's past earnings, the current demand is positively correlated with the driver's likelihood of remaining on the platform. In other words, there is evidence that the demand shock on the platform is likely to have a spillover effect on the supply side. When analyzing the effect of promotions on the focal platform, we need to take the demand-supply spillover effect into consideration. Existing marketing literature on promotions has not sufficiently considered this spillover effect, which constitute an intended contribution for our study.

Table 2.5 OLS Regression Results on Driver Retention

	Est (SE)
$\log(\text{PastWage}_{jkm,t-1})$	0.039*** (0.000)
$\log(\text{Demand}_{kmt})$	0.028*** (0.005)
Individual Fixed Effect	Yes
Market-time Fixed Effect	Yes
Observations	64,896

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; + $p < 0.1$

Combining the reduced-form evidence from Tables 2.4 and 2.5, we derive interesting correlational insights on the effects of promotions. Price discounts and service upgrades are both associated with a higher likelihood of repeated transactions for customers, i.e., positive user-stickiness effect. For a two-sided market, an increase in demand provides incentives to encourage driver retention, leading to the classic cross-side network effect. Furthermore, upgrades are positively related to a customer switching to a higher-tier in the future, which means that the spillover effect caused by upgrades goes beyond the same tier. This effect, however, is not found for price discounts: switching is only positively correlated with the past but not current discounts. Because different service tiers have different profit margins, the reduced-form evidence seems to suggest that the two promotion methods could have differential profit implications for the platform. These results, of course, are based on correlational evidence. In the next section, we develop a structural model to formally examine the role of promotions by jointly analyzing the user-stickiness effect and the spillover effect on the platform.

2.4 Model and Estimation

In this section, we present the structural models and discuss the estimation. We explicitly consider the decisions of three types of agents: consumers, drivers, and the platform. In each time period, we assume that the timing of actions is as follows. Given the current number of drivers, the platform first sets the optimal discount level and upgrade rate. After observing the discounts, consumers decide whether to make a reservation and if so, they decide which tier to book. Finally, drivers have an expectation of how much they can earn from the platform and decide whether to remain

for the period. The sequence of decisions is depicted in Figure 2.6. We first present the demand-side model followed by the supply-side model and the platform's maximization problem.

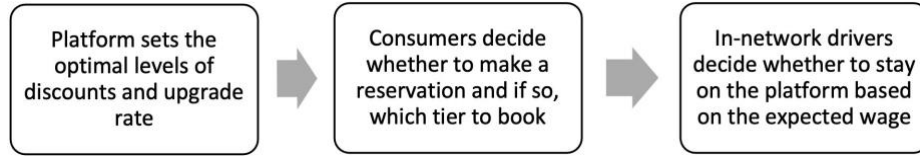


Figure 2.6 Decision Process

2.4.1 Consumer Decisions

Consumers on this platform follow a two-stage decision-making process⁵. They first decide whether to book a service from the platform, and if so, they decide which tier to choose.

In the first stage, the reservation utility for consumer i in market m at time t is specified as follows:

$$u_{imt} = \alpha_1 C_{imt} + \alpha_2 \log(\text{Dist}_{imt}) + E[\max\{U_{ikmt}\}] + \Phi_{mt}^u + \varepsilon_{imt}^u \quad (2.4)$$

In Equation 2.4, the first two terms are specific to the incidence of the focal reservation: C_{imt} is the price discount amount that i can apply to the current reservation, and Dist_{imt} is the trip distance. In our data, we can only observe the coupon usage when there is a transaction and do not observe when the coupon is distributed. Because of this, we assume that a consumer will redeem the coupon if she has one. Coupons can be applied to any tier and the trip distance is the same for a single trip. Thus, both features are captured at the first stage. Term $E[\max\{U_{ikmt}\}]$ is the inclusive value,

⁵ Hierarchical choice structures for consumer decision-making procedures have been widely adopted by previous research, such as Kannan and Wright (1991).

which captures the expected utility from choosing the best tier k for a given reservation.⁶ Term Φ_{mt}^u is the market-time fixed effect,⁷ and ε_{imt}^u is the idiosyncratic error, which is assumed to be i.i.d. after controlling for all the terms in the model. If consumer i decides not to make a reservation on the platform, s/he chooses an outside option, which could be not making a trip or choosing a better alternative. The utility for the outside option is denoted as $u_{imt}^0 = \varepsilon_{imt}^{u0}$.

Next, we specify i 's utility of choosing the service of tier k (denoted as U_{ikmt}). Tier k can take one of three values: $k = 1$ refers to the basic service; $k = 2$ is the premium service; and $k = 3$ is the luxury service. Consumer i derives the following utility when booking tier $k \in \{1, 2, 3\}$:

$$U_{ikmt} = T_k + \beta_1 S_{ikt} + \beta_2 P_{kmt} + \beta_3 \log(N_{km,t-1}) + \varepsilon_{ikmt}^U \quad (2.5)$$

Consumer i 's utility for tier k depends on the tier-specific preference, which consists of the tier fixed effect T_k and an individual's deviation from the mean. We model the systematic individual preference as a function of whether the individual has experience with tier k . Following the state-dependence specification as in Sudhir and Yang (2018), we capture past experience using S_{ikt} , which equals 1 if i had experience with k before t , and 0 otherwise. Also, the experience here is independent of the market since the service tier is based on car models, which is the same for each market. Allowing past experience to affect the current booking utility has implications in our

⁶ Here, we assume the coefficient of the expected utility equals one for tractability of the FOC calculation, which is similar to Zheng et al. (2020).

⁷ 22 % of consumers make multiple transactions. If a consumer made more than one reservation within the same month, the reservations are modeled as independent choices.

setting. Suppose the platform upgraded i 's booking to a higher tier k ($k = 2$ or 3), the consumer would receive an opportunity to experience tier k and reduce uncertainty about the service. If the experience was positive, the consumer would derive higher utility for k in the future. All else being equal, the consumer's likelihood of booking tier k increases.

Furthermore, the utility is a function of the tier-specific unit price, P_{kmt} , and the quantity of the supply N_{kmt} due to the cross-side network effect. Here N_{kmt} refers to the number of in-network drivers of tier k in market m at time t . The platform adopts a static pricing strategy so that prices are fixed and predetermined. Therefore, there is no correlation between the price and the error term which exempts the price endogeneity on the demand side. As modeling the cross-side network effect, we assume that a consumer forms an expectation of the current supply based on her past experience: the number of drivers in the past month ($N_{km,t-1}$) enters the utility to approximate the current supply size in month t . By doing so, the endogeneity due to simultaneity is accounted for. Moreover, consumers can predict the supply size based on the past waiting time. In the appendix, we show a negative relationship between the number of drivers and consumers' waiting time. More drivers on the platform will result in a quicker response time, which leads to higher utility for consumers. All else being equal, it is expected that consumers would derive higher utility in response to an increase in supply, as the higher availability of drivers can reduce the wait time and/or increase the

chance of making a successful reservation. The driver size enters Equation (2.5) in a logarithm format to reflect the diminishing return of the supply.⁸

Lastly, ϵ_{ikmt}^U is the idiosyncratic error term that follows Type I Extreme Value distribution. During the estimation, the luxury service $k=3$ is treated as the reference level, i.e., T_3 is normalized to 0. Thus, T_1 and T_2 represent the difference of customers' average preference between the luxury and the basic tier, the luxury and the premium tier, respectively. The market- and time-fixed effects are excluded from Equation (2.5) because they are constant across tiers for the same market and time combination.

2.4.2 Driver's Decision

Aforementioned, on the supply side, there are two types of drivers on the platform: in-network and out-of-network. In this section, we focus on the decisions of in-network drivers who serve the majority of the orders. In-network drivers decide whether to renew their contracts at the beginning of each month. We assume that drivers form an expectation of how much they can earn from the platform during the period. To better help in-network drivers make decisions, the platform usually provides demand and promotion data for the past six months. With historical data, drivers can estimate future demand more accurately.

If the expected earnings exceed the outside option, the driver remains with the platform for the period; otherwise, the driver leaves the platform. The driver's retention utility, denoted as v_{jkmt} , is modeled as follows:

⁸ Occasionally, rideshare platforms experience a supply shortage and need to cancel reservations causing great inconvenience and disutility for consumers. In Appendix A, we include an analysis, which confirms the negative relationship between wait time and the number of drivers for each tier.

$$v_{jkmt} = \eta_0 + \eta_1 \log[E[\tilde{d}_{kmt}] * (1 - \gamma^{in}) * P_{kmt} * \text{dist}_{mt}] + \eta_2 \text{tier} + \Phi_{mt}^v + \xi_{kmt} + \varepsilon_{jkmt}^v \quad (2.6)$$

Here, η_0 represents the average benefit of serving on this platform, which is constant across all drivers from all markets. The second term is the driver's expected wage at time t to capture the deviance from the mean, which we describe in detail next. Tier refers to the tier of driver j . Retention can vary by market and time, which is captured via the fixed effects Φ_{mt}^v . Furthermore, the retention utility is affected by unobserved shocks that are specific to the tier, market, and time, denoted as ξ_{kmt} . Lastly, ε_{jkmt}^v is the idiosyncratic error term that follows the Type I Extreme Value distribution.

A driver's expected wage is defined as the product of the expected realized demand \tilde{d}_{kmt} ⁹, the commission rate $(1 - \gamma^{in})$ paid to drivers, the unit price, and the average distance of the trips. Since the services in our setting are primarily for picking-up or dropping-off customers at airports or train stations, drivers know the average distance in each market. The unit prices and the commission rates are also known to the drivers. Therefore, the only uncertain term for drivers is future demand, \tilde{d}_{kmt} , for which the drivers would form an expectation. The expected wage enters the utility in the logarithm form to capture the diminishing return of this effect.

⁹ We model drivers' retention decision based on total demand instead of individual demand since in our context, in-network drivers are under direct control of the platform and have fixed shifts based on the contract. Also, we are focusing on how do different user growth strategies affect demand and supply. Here the direct network effect on the supply side is not our first order concern, and we leave it for future research.

Due to free upgrades on the platform, the realized demand is a function of the reservation demand d_{kmt} and the upgrading probability θ_{kmt} :

$$\begin{aligned}\tilde{d}_{1mt} &= (1 - \theta_{1mt})d_{1mt} \\ \tilde{d}_{2mt} &= \theta_{1mt}d_{1mt} + (1 - \theta_{2mt})d_{2mt} \\ \tilde{d}_{3mt} &= \theta_{2mt}d_{2mt} + d_{3mt}\end{aligned}\tag{2.7}$$

We have the following assumptions. Firstly, we assume that drivers know the distribution of the upgrade rates (θ_{kmt}) and the reservation demand (d_{kmt}). Secondly, the upgrade rate is independent of the reservation demand. Lastly, we use the true upgrade rates and the reservation demand for month t to approximate the means. These assumptions are based on the fact that the past demand patterns and promotion data are disclosed to in-network drivers. In this way, drivers can predict the future using the historical data, which validates our assumptions of using the true data to approximate the means. Alternatively, one can adopt a time-series model that predicts the demand for time t using historical data and then takes the predicted value as the mean of the demand or upgrade rate. The assumption of the independence of the two distributions is based on the low correlations between the upgrade rates and reservation demand from the data.

Assuming that drivers in each market know the distribution of the upgrade rates, the expectation is approximated based on the average upgrade rate observed in each market, leading to the following equations:

$$\begin{aligned}E[\tilde{d}_{1mt}] &= (1 - \theta_{1mt})d_{1mt} \\ E[\tilde{d}_{2mt}] &= \theta_{1mt}d_{1mt} + (1 - \theta_{2mt})d_{2mt} \\ E[\tilde{d}_{3mt}] &= \theta_{2mt}d_{2mt} + d_{3mt}\end{aligned}\tag{2.8}$$

With Equations (2.6) through (2.8), we can now derive the market share for the drivers who choose to stay on the platform. Let s_{kmt} denote the market share of the retained drivers, which can be computed as:

$$s_{kmt} = \frac{\exp(\delta_{kmt})}{1 + \exp(\delta_{kmt})} \quad (2.9)$$

where $\delta_{kmt} = v_{jkmt} - \epsilon_{jkmt}^v$.

If driver j decides to leave the platform, s/he chooses the outside option whose utility is normalized to 0. Thus, the market share for the outside options becomes:

$$s_{kmt}^0 = \frac{1}{1 + \exp(\delta_{kmt})} \quad (2.10)$$

We denote the number of retained drivers as N_{kmt} and the total number of available drivers as N_{kmt}^T . Thus, we have the relationship: $\log(s_{kmt}/s_{kmt}^0) = \log(N_{kmt}/N_{kmt}^T - N_{kmt})$. In our setting, the pool of available drivers is so large that the number of drivers who choose the outside option is much greater than the number of those who work as platform drivers, i.e., $N_{kmt}^T - N_{kmt} \gg N_{kmt}$. With the assumption that $N_{kmt}^T - N_{kmt} \approx N_{kmt}^T$, we obtain the following equations:

$$\begin{aligned} \log(N_{1mt}) &= \eta_0 + \eta_1 \log[(1 - \theta_{1mt})d_{1mt} * (1 - \gamma^{in}) * P_{1mt} * \text{dist}_{mt}] + \eta_2 \text{tier} + \eta_3 \log(N_{1mt}^T) + \Phi_{mt}^v \\ &\quad + \xi_{1mt} \\ \log(N_{2mt}) &= \eta_0 + \eta_1 \log\left[\left[\theta_{1mt}d_{1mt} + (1 - \theta_{2mt})d_{2mt}\right] * (1 - \gamma^{in}) * P_{2mt} * \text{dist}_{mt}\right] + \eta_2 \text{tier} \\ &\quad + \eta_3 \log(N_{2mt}^T) + \Phi_{mt}^v + \xi_{2mt} \\ \log(N_{3mt}) &= \eta_0 + \eta_1 \log\left[\left(\theta_{2mt}d_{2mt} + d_{1mt}\right) * (1 - \gamma^{in}) * P_{3mt} * \text{dist}_{mt}\right] + \eta_2 \text{tier} + \eta_3 \log(N_{3mt}^T) \\ &\quad + \Phi_{mt}^v + \xi_{3mt} \end{aligned} \quad (2.11)$$

Note that the retention decision is only relevant for the in-network drivers who have a month- by-month contract with the platform. In contrast, the out-of-network drivers

are not bound by contract and are recruited only when the platform experiences excessive demand. Out-of-network drivers are more costly to acquire, and their availability is less stable and more difficult to predict. Thus, the platform prefers having a stable supply of in-network drivers. The size of in-network drivers is also a key factor for the cross-side network effect. Sustaining the driver base saves potential opportunity costs for the platform, which plays an important role in the platform's marketing strategy. In the next sub-section, we specify the profit maximization problem of the platform, which formally incorporates both the choice decisions for consumers and the retention decisions for drivers.

2.4.3 Platform Decisions

As in any two-sided market, consumer utility depends on supply (i.e., the aggregated outcome of the drivers' retention decisions), and driver utility depends on demand (i.e., the aggregated outcome of the consumers' reservation and the tier choice decision). The platform can influence the cross-side dynamics via the marketing variables such as discounts and upgrades. The discount level and the upgrade rate are endogenously determined by the platform based on its profit maximization, after simultaneously considering the demand-and-supply network effects. In this section, we present our model for the platform decisions.¹⁰

¹⁰ Due to the data limitation, we only observe the coupons and upgrades when a transaction exists. Thus, we simplify the promotion decisions by assuming that the platform decides each consumer's average coupon face value and the upgrade rate for

The platform employs both in-network and out-of-network drivers and the commission rates from the two types of drivers are denoted as γ^{in} and γ^{out} , respectively. The platform's total revenue minus driver's commission equal:

$$\begin{aligned}
& P_{1mt} \text{dist}_{mt} \{ \gamma^{\text{out}} [(1 - \theta_{1mt}) d_{1mt} - q_{1mt} N_{1mt}] + \gamma^{\text{in}} q_{1mt} N_{1mt} \} \\
& + P_{2mt} \text{dist}_{mt} \{ \gamma^{\text{out}} [(1 - \theta_{2mt}) d_{2mt} - q_{2mt} N_{2mt}] + \gamma^{\text{in}} q_{2mt} N_{2mt} \} \\
& + P_{3mt} \text{dist}_{mt} \{ \gamma^{\text{out}} [(1 - \theta_{3mt}) d_{3mt} - q_{3mt} N_{3mt}] + \gamma^{\text{in}} q_{3mt} N_{3mt} \} \\
& = (\gamma^{\text{in}} - \gamma^{\text{out}}) \text{dist}_{mt} \sum_{k=1}^3 P_{kmt} q_{kmt} N_{kmt} + \gamma^{\text{out}} \text{dist}_{mt} [(1 - \theta_{1mt}) d_{1mt} P_{1mt} \\
& + (1 - \theta_{2mt}) d_{2mt} P_{2mt} + d_{3mt} P_{3mt}]
\end{aligned} \tag{2.12}$$

The total revenue is calculated as the number of orders times the per-transaction revenue, $P_{kmt} \text{dist}_{mt}$. The number of orders that are fulfilled by in-network drivers equals the number of in-network drivers times their average fulfillment. We denote the average number of orders that are fulfilled by a tier-k in-network as q_{kmt} , which is known to the platform at the beginning of each time t. Since orders are dispatched to in-network drivers first and out-of-network drivers participate only if there is excessive demand, the total number of orders that are fulfilled by out-of-network drivers equals the realized demand after free upgrades minus the orders that are served by in-network drivers. According to the platform, the commission rate from in-network drivers ($\gamma^{\text{in}} = 0.3$) is higher than that of out-of-network drivers ($\gamma^{\text{out}} = 0.2$).

each market m and time t. And our model captures the average effect of different promotions, providing rich managerial implications to platforms.

We would like to emphasize that driver retention is critical for platforms' growth, as they are competing for service providers as much as they compete for consumers. Whereas extant literature has extensively studied customer retention and customer value, in this paper, we explicitly examine the value of service providers and model how providers' retention plays a role in profit maximization for the platform. Specifically, the platform solves the market-specific optimization problem with respect to discount level (c_{mt}) and upgrade rates (from basic to premium θ_{1mt} and from premium to luxury θ_{2mt}) as follows:

$$\begin{aligned}
\max_{\{\theta_{1mt}, \theta_{2mt}, c_{mt}\}} & \sum_{k=1}^3 R_{kmt} N_{kmt} + (\gamma^{\text{in}} - \gamma^{\text{out}}) \text{dist}_{mt} \sum_{k=1}^3 P_{kmt} q_{kmt} N_{kmt} \\
& + \gamma^{\text{out}} \text{dist}_{mt} [(1 - \theta_{1mt}) d_{1mt} P_{1mt} + (1 - \theta_{2mt}) d_{2mt} P_{2mt} + d_{3mt} P_{3mt}] \\
& - c_{mt} \sum_{k=1}^3 d_{kmt} - (1 - \gamma^{\text{in}}) \text{dist}_{mt} [\theta_{1mt} d_{1mt} (P_{2mt} - P_{1mt}) \\
& + \theta_{2mt} d_{2mt} (P_{3mt} - P_{2mt})]
\end{aligned} \tag{2.13}$$

The objective function of the platform consists of three components: the benefits from retaining in-network drivers, the total revenue minus driver commissions, and the marketing cost of providing price discounts and free upgrades. In Equation 2.13, the first term represents the benefits of retaining in-network drivers for tiers 1, 2, and 3. Here, N_{kmt} represents the number of drivers for tier k who are retained on the platform in market m at time t . The marginal benefit of retaining a tier- k driver, denoted as R_{kmt} , can be understood as the opportunity cost of retaining a driver for tier k : if a driver quits, the platform has to incur acquisition costs to fill the vacancy. Intuitively, with

one more in-network driver retained, the platform saves R_{kmt} . The next term corresponds to the platform's earnings in Equation 2.12.

The remaining terms in Equation 2.13 correspond to the marketing costs of coupons and upgrades. The total cost of price discounts is the average coupon amount c_{mt} times the total demand $\sum_{k=1}^3 d_{kmt}$.¹¹ The last term corresponds to the cost of free upgrades: $(P_{2mt} - P_{1mt})\text{dist}_{mt}$ represents the total loss in revenue due to free upgrades from tier k to $k+1$.

For ease of exposition, we omit the market and time subscripts and re-write Equation 2.13 as:

$$\begin{aligned} \max_{\{\theta_1, \theta_2, c\}} \Pi = & \sum_{k=1}^3 R_k N_k + (\gamma^{\text{in}} - \gamma^{\text{out}}) \text{dist} \sum_{k=1}^3 P_k q_k N_k + \gamma^{\text{out}} \text{dist} [(1 - \theta_1) d_1 P_1 + \\ & (1 - \theta_2) d_2 P_2 + d_3 P_3] - c \sum_{k=1}^3 d_k - (1 - \gamma^{\text{in}}) \text{dist} [\theta_1 d_1 (P_2 - P_1) + \theta_2 d_2 (P_3 - P_2)] \end{aligned} \quad (2.14)$$

The FOCs of the platform's optimization problem become:

$$\begin{aligned} \frac{\partial \Pi}{\partial \theta_1} = & \sum_{k=1}^3 R_k \frac{\partial N_k}{\partial \theta_1} + (\gamma^{\text{in}} - \gamma^{\text{out}}) \text{dist} \sum_{k=1}^3 P_k q_k \frac{\partial N_k}{\partial \theta_1} - \gamma^{\text{out}} d_1 P_1 \text{dist} - (1 - \gamma^{\text{in}}) d_1 (P_2 - P_1) \text{dist} = 0 \\ \frac{\partial \Pi}{\partial \theta_2} = & \sum_{k=1}^3 R_k \frac{\partial N_k}{\partial \theta_2} + (\gamma^{\text{in}} - \gamma^{\text{out}}) \text{dist} \sum_{k=1}^3 P_k q_k \frac{\partial N_k}{\partial \theta_2} - \gamma^{\text{out}} d_2 P_2 \text{dist} - (1 - \gamma^{\text{in}}) d_2 (P_3 - P_2) \text{dist} = 0 \\ \frac{\partial \Pi}{\partial c} = & \sum_{k=1}^3 R_k \frac{\partial N_k}{\partial c} + (\gamma^{\text{in}} - \gamma^{\text{out}}) \text{dist} \sum_{k=1}^3 P_k q_k \frac{\partial N_k}{\partial c} + \gamma^{\text{out}} \text{dist} [(1 - \theta_1) P_1 d_1 \frac{\partial d_1}{\partial c} \\ & + (1 - \theta_2) P_2 \frac{\partial d_2}{\partial c} + P_3 \frac{\partial d_3}{\partial c} - c \sum_{k=1}^3 \frac{\partial d_k}{\partial c} \\ & - \sum_{k=1}^3 d_k - (1 - \gamma^{\text{in}}) \text{dist} [\theta_1 \frac{\partial d_1}{\partial c} (P_2 - P_1) + \theta_2 \frac{\partial d_2}{\partial c} (P_3 - P_2)] = 0 \end{aligned}$$

¹¹ Note that the coupon amount is independent of upgrading; thus, shifting the demand across tiers does not alter the total cost due to coupons. Therefore, the coupon cost only depends on the total demand summed over d_{kmt} .

(2.15)

With the FOCs and the estimates from the demand and supply models, we can recover the benefit parameters R_{kmt} . The derivation of FOCs is in Appendix C. So far, we have completed the model specifications for the consumers, the drivers, and the platform, respectively.

2.5 Results

In this section, we present the parameter estimates for the structural models. Based on the estimates, we conduct a counterfactual analysis that quantifies and compares the overall effects of price discounts and free upgrades.

2.5.1 Parameter Estimates

Consumer decisions

We begin by presenting the parameter estimates for consumers' reservation decisions. The results are in Table 2.6. Not surprisingly, a deeper discount is associated with a higher probability of making a reservation (0.043, $p < 0.001$). All else being equal, a consumer is more likely to book a service if she receives a higher price cut to reduce the cost of the trip. In addition, the likelihood of booking increases with the trip distance (0.360, $p < 0.001$), suggesting that the reservation service offered by the platform is more attractive when consumers need a longer ride.

Table 2.7 presents the parameter estimates for the tier-choice model. When deciding which tier to book, a consumer considers the following factors for each tier: the unit price, whether s/he has had experience with the tier, the number of drivers, and the overall preference for each tier. The price coefficient is estimated to be negative and

significant (-0.174, $p < 0.001$). The coefficient for the number of drivers is estimated to be positive and significant (0.153, $p < 0.001$), confirming a positive cross-side network effect on this platform. All else being equal, consumers derive higher utility from a larger pool of available drivers. We also find a positive and significant effect for state dependence (0.976, $p < 0.001$): the experience with a tier leads to the higher utility for that tier in the future. The positive effect of past experience could justify the economic gains for the platform offering free upgrades. Lastly, the average tier preference of the basic and premium tier is higher than that for the luxury tier (which is used as the reference level). This could indicate that on this platform, consumers may not be fully informed about the true valuation of the higher tier. Thus, providing free upgrades can help customers learn the true quality of higher tiers.

Table 2.6. Consumer Decision: Making a Reservation

	Est (SE)
Discount	0.043*** (0.000)
Log(Distance)	0.036*** (0.006)
Constant	-3.575*** (0.072)
Market-time Fixed Effects	Yes
Observations	2,960,833

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; + $p < 0.1$

Table 2.7 Consumer Decision: Tier Choice

	Est (SE)
Price	-0.174*** (0.008)
Experience	0.976*** (0.007)
Log(NumDriver)	0.153*** (0.001)
T1	1.083*** (0.023)
T2	0.830*** (0.017)
Observations	882,567

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; + $p < 0.1$

Driver decision

Next, we discuss our results on driver retention. Table 2.8 presents the estimates for Equation 2.11. After controlling for the market-time fixed effects, the expected wage is positively correlated with the number of in-network drivers (0.694, $p < 0.001$). In other words, whether or not a driver stays with the platform is positively associated with the amount of revenue that s/he expects to make on the platform at time t . This indicates that the cross-side network effect from demand to supply is also present on this platform. Furthermore, the coefficient for the tier is estimated to be negative and significant (-0.073, $p < 0.001$), indicating that drivers in lower tiers are more likely to remain on the platform than those in higher tiers.

Table 2.8 Driver Retention Decision¹²

	EST (SE)
Log(Wage)	0.694*** (0.023)
Tier	-0.073*** (0.017)
Constant	-3.286*** (0.342)
Market-time Fixed Effect	Yes
Observations	819

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$; + $p < 0.1$

Platform decision

Table 2.9 presents the estimated value of retaining an in-network driver. Overall, the value for retaining a lower-tier driver is smaller than that for retaining a higher-tier driver. This can also be understood that on this platform, the opportunity cost is higher if a higher-tier driver leaves the platform. In particular, the value for the basic-tier driver is significantly lower than that for the premium or luxury tier, as the 95% confidence

¹² Note: The estimation is conducted using the market-level share for the drivers who choose to remain on the platforms.

interval for the basic-tier driver does not overlap with those for the other two tiers. This result provides another reason the platform is incentivized to provide free upgrades: the platform gains relatively more from retaining a higher-tier driver, and thus it benefits from shifting the demand to higher-margin drivers. This payoff from upgrades is due to supply-side factors and because free upgrades can boost customer experience and thus increase future demand.

Table 2.9 The Value of Retaining Drivers¹³

	EST (SE)	95% CI
Basic (R1)	5,289.9 (3,661.7)	[4,864.5, 5,733.2]
Premium (R2)	6291.2 (2,962.8)	[5,939.7, 6,642.6]
Luxury (R3)	6,848.1 (3,452.8)	[6,438.5, 7,257.7]

2.5.2 Promotion Effectiveness

Based on the parameter estimates, we now calculate the return on investment (ROI) of different promotions. Our simulation steps are as follows. To quantify the ROI for price discounts, we keep the number of free upgrades the same as in the data but reduce the discount at time t by 50%. Given the new level of price cuts, we simulate the equilibrium consumer decisions and the subsequent in-network driver decisions at time t . Due to the cross-side network effects, we then take the updated demand and supply at time t as the input for the period $t+1$, while keeping the discount levels and upgrade rates unchanged at $t+1$. The sum of the profits for the period t and $t+1$ becomes the counterfactual profits for the reduced discounts. Thus, the difference between the

¹³ Note: The estimates are in the local currency where the platform operates. 1 local currency is equivalent to 0.16 US dollars.

baseline profit and the counterfactual profit corresponds to the value of providing deeper discounts. The ROI of discounts for the period t is then defined as the ratio between the profit gain and the cost of offering discounts. We repeat the process for all periods and calculate the average. The ROI of upgrades is calculated in a similar way.

Table 2.10 presents the ROIs for each promotion method. Overall, a one-dollar discount on this platform can bring a long-term profit of 1.853 dollars (SE = 0.313). The ROI for free upgrades is found to be higher than that for discounts because upgrades can directly impact the drivers' decisions as well as consumers' choices, whereas discounts have a direct impact only on demand but an indirect effect on supply. Between two types of updates—from basic to premium and from premium to luxury—it is *ex-ante* unclear which one has a higher return. On the one hand, the luxury-tier driver is more valuable to the platform (see Table 2.9), which would motivate the platform to offer an upgrade from premium to luxury. On the other hand, an upgrade to the premium tier can have a larger impact on demand than an upgrade to the luxury tier because consumers are more likely to choose the premium tier than the luxury tier (partially because of the larger driver pool in the premium tier). Thus, allowing consumers to experience the premium service is more rewarding in terms of bringing future demand. The net effect of these two factors would determine the relative size of the ROI for the two upgrades. Our results show that for the upgrade of basic to premium, a one-dollar investment would bring a return of 3.492 dollars (SE = 0.474), and for the upgrade of premium to luxury, the return rate is 2.591 (SE = 0.483). Interestingly, the ROI is higher when the platform upgrades a basic-tier booking to a premium tier than upgrading a premium-tier booking to a luxury tier.

Table 2.10 Return on Investment (ROI) for Promotions

	Discount (SE)	Free Upgrades	
		Basic to Premium (SE)	Premium to Luxury (SE)
ROI	1.853 (0.313)	3.492 (0.474)	2.591 (0.483)

2.6 Counterfactual Analyses

In this section, we conduct two counterfactual analyses. In the first analysis, we propose a new incentive design for drivers and discuss the impact on consumers, drivers, and the platform. In the second analysis, we unbox the mechanisms of user-growth strategies and provide insights into the different promotion effects as the platform measures. In both analyses, we explicitly incorporate the spillover effect of promotions to account for the network effect between demand and supply.

2.6.1 Providing Driver Bonuses

With a fixed marketing budget, a new platform often needs to decide how to allocate resources to provide monetary incentives for different user groups. Should the incentives be given to the customers or suppliers or split between them? In our setting, the platform provides direct incentives only to consumers in the form of coupons and upgrades. For drivers, the incentives are indirectly realized through increased wages (as a result of increased demand). To better understand this policy, we examine a new incentive design in which the platform considers providing monetary incentives directly to the drivers directly, such as giving drivers bonuses if they continue providing service on the platform. This type of incentive has been adopted by on-demand platforms such as Uber, Lyft, and Instacart with a clear intention to attract workers.

With this new policy, in addition to the regular revenue earned from taking orders, drivers can receive a bonus if they decide to remain on the platform. In our calculation, we set the bonus budget at 60% of the marketing cost of price discounts¹⁴ and evenly distribute the bonuses among active drivers at the end of each period. With this retention bonus, drivers are more likely to remain on the platform. The incremental driver retention will lead to higher demand and greater platform revenue through the chain of network effects. To better illustrate the effect of this bonus policy, we do not change the promotion levels. The results are presented in Table 2.11.

Our calculation shows a positive ROI for this bonus policy. The number of in-network basic, premium and luxury drivers increase by 6.66%, 5.08%, and 12.98%, respectively. Specifically, as the pool of in-network drivers increases, more orders will be fulfilled by these drivers. In Figure 2.7, we compare the updated fulfillment rate by in-network drivers with the actual rate and observe an increase for all three tiers. More orders fulfilled by in-network drivers has two benefits: consumers incur shorter wait times when booking because in-network drivers provide a more stable supply than out-of-network counterparts; and the platform earns a higher commission due to the rate difference between the two types of drivers. Taking both the revenue and the bonus budget into consideration, the policy increases platform profits by 4.08%. This result is consistent with Ming et al. (2019): they show that increasing drivers' commission rate leads to a profit increase. The ROI of the bonus policy is 1.662, suggesting that providing a bonus to in-network drivers can be a potentially viable user-growth strategy.

¹⁴ 60% of the marketing cost of price discounts is equivalent to around 10% of in-network drivers' monthly wage. We also provide robustness check for alternative percentages in Appendix B.

The ROI for driver bonus is also lower than that for coupons and free updates, justifying the platform’s current practice: with a fixed budget, the marketing money is allocated to promotion methods with a higher ROI.

Table 2.11 Provide Bonus to In-network Drivers

	Basic (SE)	Premium (SE)	Luxury (SE)
Demand	0.36% (0.000)	0.32% (0.000)	0.36% (0.000)
Supply	6.66% (0.018)	5.08% (0.011)	12.98 (0.033)
Profit		4.08% (0.014)	
ROI		1.662 (0.232)	

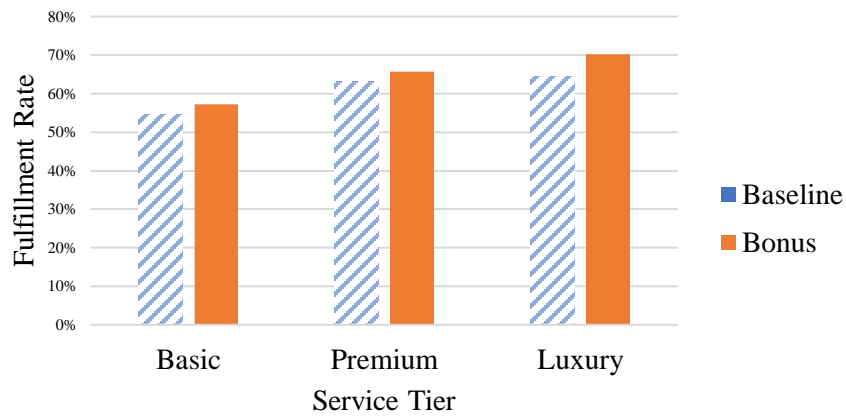


Figure 2.7 Fulfillment Rate by In-network Drivers

2.6.2 Marketing Mix at Different Stages

Our second analysis examines how the platform should adjust the marketing mix of discounts and upgrades at different stages of its life cycle. As the platform penetrates the market, the percent of experienced consumers would increase and the percent of new customers would decrease, affecting the relative effectiveness of discounts versus upgrades. With more experienced customers, naturally, the demand-side “sampling” effect of free upgrades would be curtailed. However, the spillover effect on the driver

side is still present. As we show in Section 2.6.1, having more premium- and luxury-tier drivers reduces the wait time for consumers and increases the commission revenue for the platform. As for the discounts, interestingly, our model shows that coupons are more effective for new customers than for experienced customers because of the positive coefficient for experience, as shown in Equation 2.5. However, because the effectiveness of upgrades is *ex-ante* unclear, it is also unclear how the equilibrium levels of discounts and upgrades would change in response to more experienced customers.

To conduct this analysis, we compare the different stages of the platform when the percentage of consumers who have used the basic and premium tier services increases from 5% to 40%. We compute the equilibrium discounts and upgrades, and calculate the updated platform profits and the ROI of the two promotions. Figure 2.8 shows the changes in ROI at different stages. As the percent of experienced customers increases, the ROI of upgrades increases whereas the ROI of coupons decreases. Again, this is caused by two forces: (1) coupons are less effective for experienced customers on this platform, and (2) the positive effect of upgrades on the supply side dominates the diminishing return on the demand size. Interestingly, we also note that the importance of upgrading to different tiers matters in this analysis. In the early stages, upgrades from the basic to the premium tier are more efficient than those from the premium to luxury tier. However, as the platform penetrates the market, upgrades to the luxury tier increase the ROI much faster than upgrades to the premium tier because new consumers tend to start with lower tiers and the importance of sampling a higher tier naturally grows over time. Our findings indicate that when 40% of the consumers experience the

basic and premium tiers, the effectiveness of the two types of upgrades becomes similar. This trend provides managerial implications for the platform: discounts are more effective in the early stage, and upgrades become more effective as the platform gains greater market share. Specifically, the platform should consider leveraging the investment on upgrades between the tiers at different stages. Early on, the platform should focus on upgrades to the premium tier, but as it expands, the platform should consider providing more upgrades from the premium to the luxury tier.

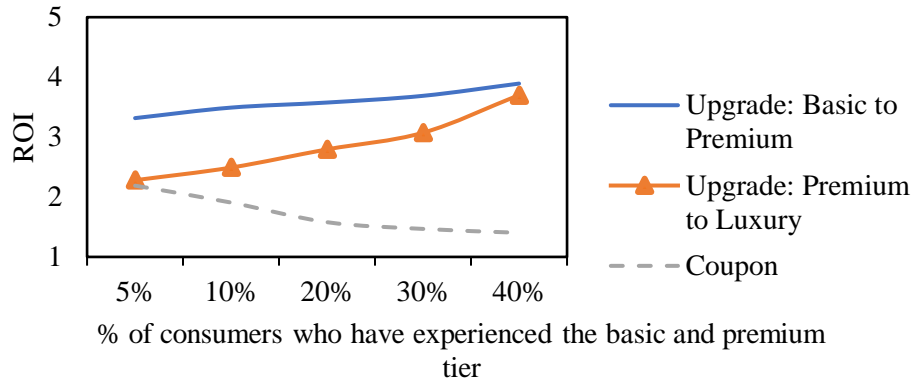


Figure 2.8 Changes of ROI

2.7 Conclusion

This paper evaluates the effectiveness of different user-growth strategies on a two-sided platform. Using transaction-level data from a leading transportation service platform, we first show that two promotion strategies, price discounts and free upgrades, help increase consumer stickiness and make drivers more likely to remain on the platform. Moreover, given a free upgrade, a consumer is more likely to switch to a higher service tier for the next booking, but this switching effect is not found for price discounts. Our structural model takes a utility-based approach to model the consumers',

drivers' and platform's decisions. We find that the benefit of retaining a higher-tier driver is larger than a lower-tier driver. Combining the spillover effects on the demand and supply side, free upgrades are more efficient than price discounts in terms of ROI. In our counterfactual analyses, we find that providing incentives directly to drivers can attract more drivers and benefit the platform. A platform should also adjust its optimal marketing mix of user-growth strategies in different stages. Our second analysis indicates that the relative importance of discounts drops while the importance of upgrades rises as the platform matures.

This paper suffers from a few limitations. First, we assume that consumers are myopic. With multiple experiences with the platform, one may argue that consumers may learn the upgrade probability gradually and form an expectation of the upgrade rate. However, due to the data limitation, we are unable to track and estimate consumers learning behavior. Our data covers 21 months, and due to the specific service type (airport transfers), only 22% of the consumers have multiple transactions. And over 80% of those with multiple transactions only booked two or three times. Learning the upgrades via transactions will make consumers more strategic in choosing which tier to book. Thus, it is important for the platform to examine consumers' strategic behavior and adjust the upgrade rates dynamically.

Secondly, we assume that drivers' retention decision is based on the total demand and we do not consider direct network effects on the supply side. To model driver competition, a more comprehensive model replaces the total demand d_{kmt} with d_{kmt}/N_{kmt} as calculating the expectation. In this way, drivers decide whether to

renew contracts based on their own expected wage instead of the total salary for all drivers.

Third, when modeling the platform's decision, we build a reduced-form dynamic model: the platform decides the optimal promotion levels by maximizing its profit during current time period t . We include the term R_{kmt} to account for the discounted value of retaining a tier- k driver in the future. However, in a full-fledged dynamic model, the platform determines the promotion levels in long run: current decisions will influence the future revenue. Thus, instead of only considering the current revenue for one time period, the platform maximizes the expected revenue in the future. We hope future research can address these limitations.

Chapter 3: Hybrid Employment Modes for Ride-hailing Systems

3.1 Introduction

The original business model of ride-hailing services such as Uber or Lyft positioned the service provider as an intermediary that matched independent contractor ride-providers with potential customers. Such businesses formed a fundamental component of the "gig economy" that enabled a workforce largely made up of personnel who, as independent contractors, were responsible for funding basic life-support services traditionally provided to regular employees. Over time a debate emerged as to whether to classify independent-contract drivers as employees in gig companies and, as such, to provide them a benefits package: In 2019, California passed a bill that classifies gig workers as employees instead of independent contractors. This legislation, which is called California Assembly Bill 5 or AB5, requires gig companies to adhere to minimum wage laws and provide benefits and protections to their workers. However, supported by Uber, Lyft, DoorDash, and Instacart, California's Proposition 22 was approved in 2020, stipulating that ride-hail and food delivery workers could be independent contractors. In 2021, a state judge in California ruled Proposition 22 to be unconstitutional by stating that "It appears only to protect the economic interest of the network companies in having a divided, ununionized workforce, which is not a stated goal of the legislation." In 2022, Washington State lawmakers passed the bill, HB 2076, which gives ride-hail workers new benefits—sick pay, a process to appeal deactivations, protections against retaliation, and workers' compensation (Feliz Leon 2022).

These new regulations affect both gig companies and their workers. Once those workers are classified as employees, gig companies have to provide benefits such as health insurance and unemployment insurance to their employees, which will increase the cost and significantly affect the profit margin. It can further impact consumers since some of the costs might be passed to the users if the platforms want to preserve their profit margin. These policies influence workers as well. Being classified as employees will guarantee the minimum wage and provide greater protections for drivers. However, some drivers are against it: they usually work part-time for those platforms and value the flexibility and freedom as an independent contractor. Such workers value the ability to choose when and how long to work for themselves instead of following fixed schedules.

Therefore, how to balance flexibility and benefits for drivers with the profitability for platforms is a fundamental question for the future of these industries. Prior research studied the cases when drivers are self-scheduling, and on-demand platforms hire independent contractors only. On the supply side, literature investigated driver behavior under extensive margin and intensive margin, i.e., whether to participate in the platform and, if so, how many hours to work. For example, Banerjee et al. (2015), Bai et al. (2018), Taylor (2018), and Gurvich et al. (2017) study drivers' participation decisions in on-demand services. Baron (2018), Benjaafar et al. (2019), and Bimpikis et al. (2019) discuss both the extensive and intensive margins. Cachon et al. (2017) and Yu et al. (2019) consider a two-stage decision for drivers: at the first stage, drivers decide whether to join the platform, and at the second stage, drivers decide whether to work during a particular time. From the platform's perspective, pricing strategies are

crucial for its profitability. Static pricing and dynamic pricing also called surge pricing, are widely adopted by ride-hailing platforms. Castillo et al. (2017) shows that surge pricing can prevent the inefficiency in matching the demand and supply to avoid a "wild goose chase." Cachon et al. (2017) argues that surge pricing generally achieves nearly the optimal profit, and as labor becomes more expensive, all stakeholders can benefit from surge pricing. By contrast, Banerjee et al. (2015) and Chen and Hu (2019) show that dynamic pricing does not necessarily yield higher performance than static pricing. However, facing new challenges, previous business models may not be applicable. Policies such as AB5 and Proposition 22 feed a growing literature that explore hybrid employment modes: platforms hire both employees and independent contractors (Dong and Ibrahim 2018, Hagiú and Wright 2019, Chakravarty 2021, Lobel et al. 2021). For example, Chakravarty 2021 shows that a blended platform capacity becomes viable if the wage rate is moderate, the pool of independent drivers is large, and the ride-seeker market is large. However, only Hagiú and Wright (2019) models the benefits paid to the employees, which is also part of our focus.

In this chapter, we model an on-demand platform that adopts a hybrid employment mode. In our baseline model, drivers have the flexibility to decide how long to work and whether to be an employee or a freelancer. Those who work over a certain number of hours will be classified as employees and receive a benefits package. The platform is a profit-maximizer and decides the optimal price based on the required benefit amount. As the benefit amount increases, the platform's profit decreases, which explains why gig companies spent over \$200 M supporting Proposition 22 to allow business models with drivers as independent contractors. Moreover, we show that

higher benefits make consumers and full-time drivers better off but decrease part-time drivers' welfare as well as overall social welfare. We further analyze alternate policies that balance the platform's profitability and drivers' welfare. We consider the case that the policymaker requires that a minimum portion of the workforce consist of full-time employees. In this case, the platform incurs less profit loss and has to employ a certain number of employees. Full-time drivers' welfare decreases but part-time drivers are better off.

The remaining of the paper proceeds as follows. In Section 3.2, we review relevant literature and discuss our contributions. We specify our analytical model in Section 3.3 and provide numerical studies in Section 3.4. Section 3.5 concludes.

3.2 Literature Review

The emergence and rapid growth of on-demand service platforms has attracted extensive research studies. Overviews are provided by Benjaafar and Hu (2020), which summarizes sharing economy business models and Wang and Yang (2019), which proposes a general framework to describe rides-sourcing systems specifically. Our research is related to three streams of literature: platform pricing decisions, supply models, and government regulations.

Pricing for on-demand platforms is distinguished from pricing for traditional service systems by the following features: i) capacity affects demand, and vice versa; ii) capacity can be controlled only indirectly via wages and prices; iii) capacity and demand vary temporally and spatially (Benjaafar and Hu 2019). Setting prices, wages, and commissions thus not only impact profit margins but also are fundamental to

control of the underlying system. Considerable literature discusses platforms' pricing decisions, including questions about how much a platform should charge consumers and pay its providers. Hu and Zhou (2020) shows that the optimization problem that jointly considers prices and wages can be reduced to a one-dimensional problem; they also show that the platform's optimal price has a U-shape relationship with the required minimum wage. Bai et al. (2018) shows that it is optimal for the platform to charge a higher price when demand increases. However, the optimal price is not necessarily monotonic when provider capacity or the waiting cost increases. This result is consistent with Benjaafar et al. (2019), which shows the optimal price at first increases with labor pool size and then decreases. In addition to pool size and compensation, Gurvich et al. (2017) allows the firm to impose a cap on the number of active agents in a period. They solve a multiperiod newsvendor problem with self-scheduling capacity under those three control levers. Taylor (2018) examines how consumers delay sensitivity and providers' independence impact the platform's price and wage decisions in queueing system.

Several papers compare dynamic pricing with static pricing. Cachon et al. (2017) finds that the optimal dynamic contract can substantially increase the platform's profit compared with contracts with a fixed price or fixed wage. They also numerically show that a commission contract is near-optimal. Hu and Zhou (2019) studies the performance of flat commission contracts, under which the wage is proportional to price regardless of what price is charged. They show that the platform achieves at least 75% of the optimal expected profit by using an optimal flat-commission contract. Banerjee et al. (2015) indicates that dynamic pricing does not necessarily yield higher

performance than static pricing while dynamic pricing is more robust to fluctuations in system parameters compared to static pricing.

For supply models in traditional labor economics, income rates influence labor supply behavior by affecting intensive and extensive margins. Here intensive margin refers to the number of working hours per driver, while the extensive margin refers to the number of drivers participating in the labor market (Cahuc et al. 2014). On-demand service providers decide whether to work and if so, when and how long to work based on the characteristics of the compensation package. Extensive margin (the participation decision) has been greatly discussed in the on-demand service literature. For example, Banerjee et al. (2015), Bai et al. (2018), Taylor (2018), and Gurvich et al. (2017) assume that a driver will participate only if the expected wage is higher than a certain reference wage. In this chapter we model both the extensive and intensive margins: drivers decide whether to work for the platform and how long to work. Similar to Baron (2018) and Benjaafar et al. (2021), drivers trade off positive utility from wages with negative utility from devoting time to working on the platform (opportunity cost).

Another aspect of our research involves comparing different workforce models. On-demand businesses like Uber and Lyft solely hire independent contractors: the platforms benefit from a cheaper workforce and agents enjoy their flexibility. Existing literature has extensively explored the system with independent contractors only (Cachon et al. 2017, Gurvich et al. 2017, Taylor 2018).

However, the on-going public debate on whether to classify drivers as employees demonstrates the need for research on blended workforce models. Dong and Ibrahim (2020) uses a queueing-theoretic framework to study cost-minimizing staffing

decisions in service systems with a blended workforce that includes both full-time and part-time workers. Chakravarty (2021) and Lobel et al. (2021) adopts a newsvendor style model to study the platform's maximization problem. However, in their analysis, employees are homogeneous and considered as a fixed cost. Our model investigates the pricing and staffing decisions for the platform where all potential drivers make employment status decisions (full- vs. part-time) and decide on the number of hours worked based on the compensation package. Our study will provide insights on how on-demand platforms should respond to new regulations and drivers' behavior changes.

3.3 Model

We consider an on-demand ride-hailing platform such as Uber or Lyft. We start with a relatively simple base model and then, through sensitivity analysis and extensions model more complex platform and driver options. We assume uniform orders (ride lengths) with each consuming one hour, which includes ride time and time spent driving to start of trip. For the demand side, we assume that the platform faces a linear demand curve,

$$D = a - kp \tag{3.1}$$

This is consistent with Cachon et al. (2017) and should be a good approximation over the limited range of realistic prices. Here p is the average price per order and a represents the total demand. k represents the slope of the demand function. For the supply side, we assume that the total available driver pool size is N . Drivers decide how long to work based on their opportunity costs and the revenue share from the platform. We denote the commission rate that the platform pays to drivers as δ . That

is, a driver is paid δp for serving each order. For a particular demand realization, price, and commissions, it is possible that the total supply exceeds demand. When drivers are not fully utilized, we denote the driver utilization rate, i.e., the probability of a driver is not idle, as ϕ . Then, a driver who works h hours earns $\phi\delta p h$. We employ a driver opportunity cost function based on total driver working hours h : $ch + \frac{1}{2}\theta h^2$. The opportunity cost contains two parts: a linear component and a quadratic component. The linear component describes the average per hour cost for a driver. Here we assume that the linear cost is heterogeneous among drivers and $c \sim U(0, \bar{c})$. Previous literature shows that as a driver works longer, s/he will be less willing to work, which is equivalent to an increasing marginal cost. Thus, the quadratic form describes drivers' aversion to working longer. θ is a parameter to quantify the degree of aversion.

On-demand platforms currently hire independent contractors only who are not provided a benefits package. However, under the proposed regulations, platforms may be required to provide benefits to drivers who work over a certain minimum number of hours. That is, such drivers effectively become employees and not contractors. This section will first discuss the business model with contractors only and then with a blended workforce. When the platform adopts a hybrid employment mode, we consider two scenarios: (1) drivers who work over the threshold are automatically classified as employees, i.e., Driver-controlled or DC model; (2) the platform chooses who to hire as employees and sets a cap for the number of employees, i.e., Platform-controlled or PC model.

The proofs of all the results are provided in the appendix.

3.3.1 Contractors Only: Baseline

Under the contractor-only model, drivers decide on their level of working hours by trading off wages earned and opportunity cost. A driver's utility of working h hours is $U(h) = \phi\delta ph - ch - \frac{1}{2}\theta h^2$. Thus, the optimal working hours $h^*(c)$ for a driver with the linear cost c equals $\max\left\{\frac{\phi\delta p - c}{\theta}, 0\right\}$. Here we can observe that drivers will work for the platform only if $c \leq \phi\delta p$. Also, as c decreases, drivers are willing to work longer. Knowing how long each driver works and the driver pool size N , we can calculate the platform's total capacity:

$$S = N \int_0^{\phi\delta p} \frac{h^*(c)}{\bar{c}} dc = \frac{N(\phi\delta p)^2}{2\theta\bar{c}} \quad (3.2)$$

This expression uses the assumption that each productive hour worked executes exactly one job. Since the total demand $D = a - kp$, the realized demand equals the minimum of the demand and the platform capacity, i.e., $\min\{D, S\}$ and the platform's profit is: $(1 - \delta)p \min\{D, S\}$. To solve the platform's profit maximization problem, we first observe that in any optimal solution we have that supply is at least as large as demand, i.e., the realized demand is always equal to $a - kp$.

Lemma 1 *An optimal solution always yields a realized demand $a - kp$, i.e., demand is less or equal to supply.*

Lemma 1 indicates that we have only two scenarios: i) at the optimal price demand falls short of the supply and driver utilization is less than one, which effectively rations demand among drivers; ii) the price is set so that demand exactly matches the supply. Knowing that realized demand is $a - kp$, we can then write down the platform's profit optimization problem as BOPT:

$$\max_{p, \phi} \Pi = (1 - \delta)(a - kp)p$$

$$s. t. \quad a - kp \leq S$$

$$\phi = \frac{D}{S}$$

$$(3.1), (3.2)$$

In problem BOPT, we have two decision variables: price per order and driver utilization rate. The first constraint follows from Lemma 1, and the second equation enforces the relationship between the driver utilization rate, demand and supply. By solving problem BOPT, we have:

Proposition 1 The optimal price for BOPT, is defined by $p^* = \max\{p^c, p^m\}$, where

$p^c = \frac{a}{2k}$ is the unconstrained optimal price and $p^m = \frac{\sqrt{(k\theta\bar{c})^2 + 2N\theta\delta^2 a\bar{c}} - k\theta\bar{c}}{N\delta^2}$ is the

market-clearing price when $\phi = 1$.

3.3.2 Blended Workforce: DC Model

Now we consider the situation that platforms hire both employees and contractors. Employees must work at least h_0 hours and they receive a fixed benefit amount, which is denoted as B . Then a driver's utility of working h hours becomes: $U(h) = \phi\delta ph - ch - \frac{1}{2}\theta h^2 + B\mathbb{1}\{h \geq h_0\}$. Again, each driver's decision is based on their particular opportunity cost value c . There are two types of employees: drivers who have lower opportunity costs and will work over h_0 regardless of the benefits; those with slightly higher opportunity costs who are willing to work h_0 based on the extra value provided by the benefits. Drivers with relatively high c will work part-time and those with the highest opportunity cost will not work for the platform.

There are different scenarios which depend on the price and benefits package. For example, the platform may hire independent contractors only if the required benefit amount is too high. In this essay, our focus is the mixed employment mode. Thus, here we only discuss the scenario when the platform hires both contractors and employees, which requires $\phi\delta p > \theta h_0 > \sqrt{2B\theta}$.

In summary, we have the following proposition which determines driver categories from the overall pool:

Proposition 2

For the DC Model, assuming that $\phi\delta p > \theta h_0 > \sqrt{2B\theta}$,

1. Drivers with $c \in (0, \phi\delta p - \theta h_0]$ work full-time regardless of the benefits;
2. Drivers with $c \in (\phi\delta p - \theta h_0, \phi\delta p - \theta h_0 + \sqrt{2B\theta}]$ are the switchers who work as employees when benefits are provided;
3. Drivers with $c \in (\phi\delta p - \theta h_0 + \sqrt{2B\theta}, \phi\delta p]$ will work part-time.
4. Drivers with $c \in (\phi\delta p, \bar{c}]$ will not work for the platform.

From proposition 2, the number of drivers who always work full-time equals $\frac{N(\phi\delta p - \theta h_0)}{\bar{c}}$ which is larger than 0. The number of the switchers equals $\frac{N\sqrt{2B\theta}}{\bar{c}}$. The number of part-time drivers is $\frac{N(\theta h_0 - \sqrt{2B\theta})}{\bar{c}}$ which is also larger than zero. We also describe how drivers make their decisions based on their opportunity costs in Figure 3.1. Drivers with $c \in (0, \phi\delta p - \theta h_0]$ work $h^*(c)$ hours per week. The switchers work h_0 hours. Thus, employees work $N \cdot E[h^*(c) \mathbb{1}\{c \in (0, \phi\delta p - \theta h_0]\}] + N \cdot E[h_0 \mathbb{1}\{c \in (\phi\delta p - \theta h_0, \phi\delta p - \theta h_0 + \sqrt{2B\theta}]\}]$ hours in total. Each contractor works $h^*(c)$

hours, so all part-time drivers contribute $N \cdot E[h^*(c) \mathbb{1}\{c \in (\phi\delta p - \theta h_0 + \sqrt{2B\theta}, \phi\delta p]\}]$ hours. Summing up all working hours, the platform's capacity becomes

$$S = \frac{N(\phi\delta p)^2}{2\theta\bar{c}} + \frac{NB}{\bar{c}} \quad (3.3)$$

See appendix for the derivation of S . Also, we find that by providing benefits, more drivers are incentivized to work longer.




Figure 3.1 Blended Workforce: DC Model

Similarly, we can show that the optimal solution always yields a realized demand $a - kp$, i.e., demand is less or equal to supply. Then the platform solves the following optimization problem DCOPT:

$$\max_{p, \phi} \Pi = (1 - \delta)p(a - kp) - \frac{NB(\phi\delta p - \theta h_0 + \sqrt{2B\theta})}{B\bar{c}}$$

$$s. t. \quad a - kp \leq S$$

$$\phi = \frac{D}{S}$$

$$(3.1), (3.3)$$

By solving problem DCOPT, we have:

Proposition 3 When $\delta \leq 0.8$, the optimal price for DCOPT is given by $p^* = \max\{p^c, p^m\}$, where $p^c = \frac{a}{2k}$ is the unconstrained optimal price and $p^m =$

$\frac{\sqrt{(k\theta\bar{c})^2 + 2N\theta\delta^2(a\bar{c} - NB)} - k\theta\bar{c}}{N\delta^2}$ is the market-clearing price when $\phi = 1$.

The platform chooses between the unconstrained optimal price and the equilibrium price. The optimal price depends on the total demand and the slope of the demand function. The equilibrium price is a decreasing function of the benefit amount. With higher benefits, more drivers are willing to work longer to obtain the benefits and the platform capacity increases, motivating the platform to lower the price to induce more demand. However, if the benefit amount is too high, the equilibrium price is less than p^c , and the optimal price will be independent of B .

Similar to Cachon et al. (2017), we measure the consumer surplus using $CS = (a - kp)^2/2$. Drivers' welfare is denoted as DW , which equals their total wage minus the opportunity cost. Social welfare is the sum of consumer surplus, drivers' welfare and the platform's profit, i.e., $SW = CS + DW + \Pi$. The calculation details are shown in the appendix.

3.3.3 Blended Workforce: PC Model

One key factor in the debate on whether to classify drivers as employees is the profit margin of the platforms. In the DC model, the platform provides benefits to all drivers whose working hours exceed h_0 . This gives the platform minimal control over the size of the employee pool and could lead to a huge profit loss. For example, the platform's profit drops around 40% when $B=150$. To leverage the profitability of the platform and drivers' benefits, we consider an alternative model that allows the platform to hire a limited number of employees while guaranteeing that a minimum proportion of all work will be fulfilled by employees. In this way, the profit loss due to the required benefits is capped, making this alternative policy a viable solution from the platform's perspective.

The cap of the number of employees hired by the platform is denoted as \bar{n} , then the platform only pays at most $\bar{n}B$ for the benefits package. The platform can choose who to hire as employees. Here we have the following assumption:

Assumption 1: The platform chooses the most “enthusiastic” drivers as employees, i.e., those with the lowest opportunity cost c .

Like the DC model, we have two types of (potential) full-time drivers: those who will work h_0 hours regardless of the benefits and the switchers, who will work full-time only when benefits are provided. The number of the former type is denoted as n_{f_1} and the number of the latter is n_{f_s} . Here we assume that $n_{f_1} \leq \bar{n}$. From Section 3.2, we know that $n_{f_1} = \frac{N(\phi\delta p - \theta h_0)}{\bar{c}}$ and $n_{f_s} = \frac{N\sqrt{2B\theta}}{\bar{c}}$. However, not all the switchers will be hired as employees by the platform and paid benefits. When $n_{f_1} + n_{f_s} \geq \bar{n}$, only some of the potential switchers are paid benefits. Based on assumption 1, the platform chooses the potential switchers with the lowest c as employees. If we denote the fraction of potential switchers who will be hired as employees as t , then we have $t = \min\left\{\frac{\bar{n} - n_{f_1}}{n_{f_s}}, 1\right\}$.

Proposition 4

For the PC Model,

1. Drivers with $c \in (0, \phi\delta p - \theta h_0]$ work full-time regardless of the benefits;
2. Drivers with $c \in (\phi\delta p - \theta h_0, \phi\delta p - \theta h_0 + t\sqrt{2B\theta}]$ are the potential switchers who work as employees when benefits are provided;
3. Drivers with $c \in (\phi\delta p - \theta h_0 + t\sqrt{2B\theta}, \phi\delta p]$ will work part-time.
4. Drivers with $c \in (\phi\delta p, \bar{c}]$ will not work for the platform.

We illustrate the driver categories based on their opportunity costs in Figure 3.2. Similar to the DC Model, we have four types of drivers: drivers who always work full-time regardless of the benefits, the switchers who work full-time because of the benefits, drivers who always work part-time and those who are not working for the platform. Drivers with $c \in (0, \phi\delta p - \theta h_0]$ work $h^*(c) = \frac{\phi\delta p - c}{\theta}$ hours per week. The switchers work h_0 hours. Under this policy, not all potential switchers receive the benefits and become employees. We have fewer switchers than the baseline model, and with t the fraction of potential switchers who become employees. Thus, the total working hours of employees are equal to $N \cdot E[h^*(c)\mathbb{1}\{c \in (0, \phi\delta p - \theta h_0]\}] + N \cdot E[h_0\mathbb{1}\{c \in (\phi\delta p - \theta h_0, \phi\delta p - \theta h_0 + t\sqrt{2B\theta}]\}]$. Each contractor works $h^*(c)$ hours, so all part-time drivers contribute $N \cdot E[h^*(c)\mathbb{1}\{c \in (\phi\delta p - \theta h_0 + t\sqrt{2B\theta}, \phi\delta p]\}]$ hours. Summing up all the working hours and simplifying the platform's capacity, we have

$$S = \frac{N(\phi\delta p)^2}{2\theta\bar{c}} + \frac{NBt^2}{\bar{c}} \quad (3.4)$$

See the appendix for the derivation of this equation. The new optimization model POPT is obtained by using (3.4) instead of (3.3) in BOPT and we solve POPT numerically in our experiments.

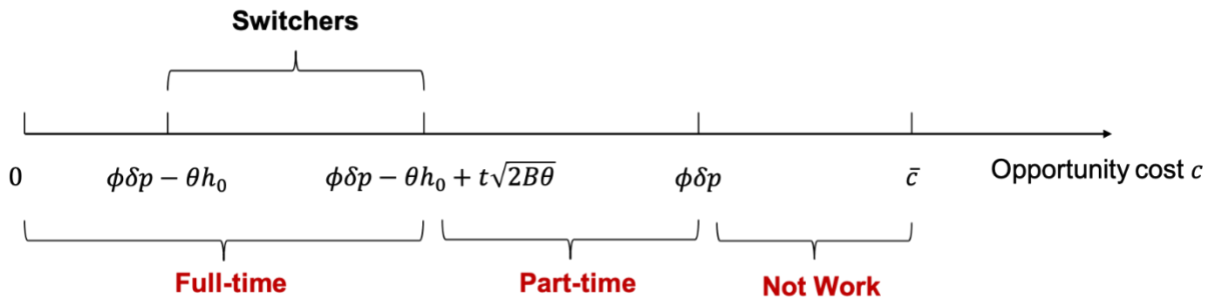


Figure 3.2 Blended Workforce: PC Model

3.4 Numerical Study

In this section, using the models just derived, we study the influence of the new regulations and alternative policies on the platform, drivers and consumers. The public debate and proposed regulations involve somehow inducing ride-hailing platforms to classify some or all drivers as employees who would receive a benefits package. Our baseline model represents the current driver-contractor model used by the major platforms. The DC and PC models both support a blended workforce that includes both employees and contractors. If a move toward supporting employees is made a blended workforce approach seems eminently reasonable as some part of the workforce would very likely wish to retain the contractor flexibility and lower working hours. We first compare the baseline model with the DC model where drivers decide how long to work and those who work over h_0 hours are classified as employees and receive the benefit amount B . This approach seems to be the implicit assumption underlying the proposed regulations. However, as we will show this approach can lead to drastic profit reductions for the platforms and even non-viable business models. As an alternative, the PC models allows the platform to control the mix of full-time employees and contractors. This approach mitigates the platform profit loss. In our analysis we analyze how different policies affect impact the welfare of the platform, drivers and consumers.

3.4.1 Baseline vs. DC Model

Proposition 22 requires that gig companies provide partial health care subsidies to drivers who average at least 25 hours per week of “engaged time,” or the periods when they are en route to pick up passengers and when they are transporting passengers in the vehicle. (Manzo et al. 2022). In this way, we use 25 hours as the threshold (h_0) to

determine whether a driver is classified as an employee. According to Uber, drivers keep 75 percent of the base fare, which means δ equals 0.75. Drivers' opportunity cost is a non-linear function of working hours h : $ch + \frac{1}{2}\theta h^2$. c is uniformly distributed on $[0, \bar{c}]$. θ quantifies the degree of drivers' aversion to working longer; based on other uses in the open literature and some initial experiments, we set θ equal to 1. Manzo et al. (2022) analyzes data on more than 22 million rideshare trips in Chicago from January 2021 through June 2021 and calculates drivers' hourly gross income, which is equal to \$27. So we adjust \bar{c} and parameters in the demand function (a, k) to coordinate with the real-world data.

Figure 3.3 shows how different indicators vary with the weekly benefit amount B , including the price, driver utilization rate, profit, the number of drivers, drivers' welfare, and social welfare. From Figure 3.3, we can observe that as B increases, the platform first chooses the equilibrium price, which is a decreasing function of B , then adopts the unconstrained optimal price which is independent of B . The average cost per trip decreases from \$42 to \$40. The platform's profit decreases with B , which is consistent with strong gig company opposition to providing benefits. For example, the profit margin drops 60% when B is 150 (about 20% of a full-time driver's weekly wage). Benefits motivate drivers to work longer: with higher benefits, more drivers are willing to be employees and full-time drivers' welfare increases. However, part-time drivers are hurt by higher benefits: their welfare decreases with the benefit amount because of the reduced hourly wage. In total, drivers' welfare first falls and then increases. And if the benefit is high enough, all contractors would like to work full-time (although this represents an extreme / unrealistic case where platform profits have

become negative). Social welfare decreases with B since the profit loss and the reduced welfare for contractors dominate the increase of consumer surplus as well as employees' welfare.

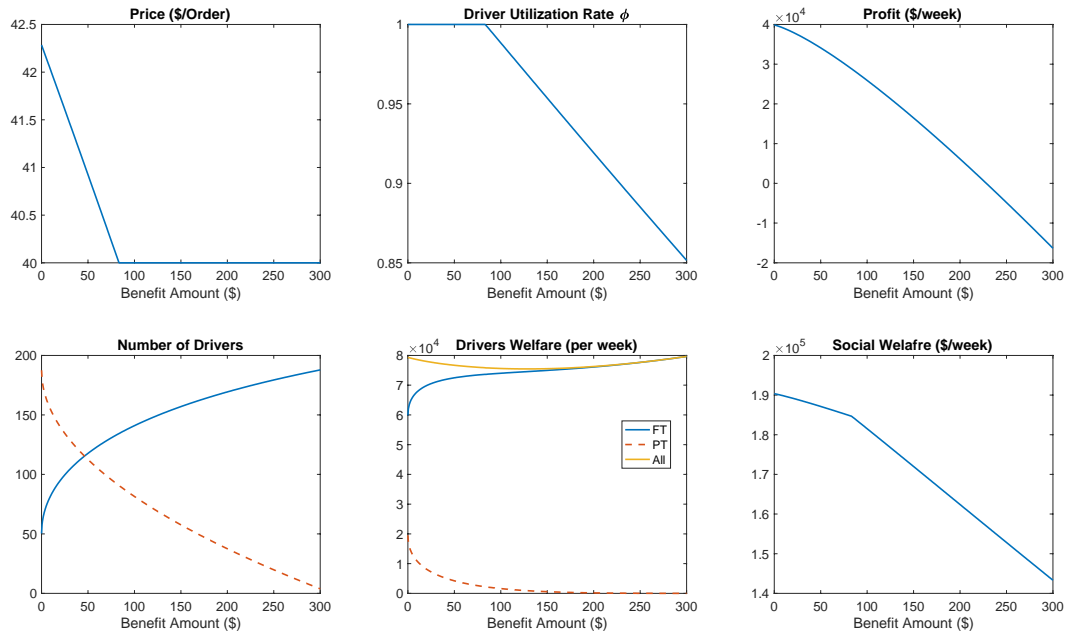


Figure 3.3 Baseline vs. DC Model

3.4.2 DC vs. PC Model

Figure 3.4, Figure 3.5, and Figure 3.6 compare the baseline policy with the new policy numerically. Particularly, Figure 3.4 shows the changes in the price, driver utilization rate, profit, consumer surplus, drivers' welfare, social welfare, full-time and part-time drivers' welfare with B; Figure 3.5 indicates how the number of drivers and drivers' working hours are affected by the benefits; Figure 3.6 compares drivers' welfare for three categories: drivers who always work full-time, potential switchers and those who always work part-time.

Here we use $\bar{n} = 100$ as the cap for the number of employees, which is around 17% of the total available drivers. While this might seem low, as we will show later, a much higher percent of the orders are fulfilled by employee-drivers. We first look at the price. In both cases, the price first decreases and then remains the same. However, the price is higher under the new policy than the baseline. The platform benefits from the new policy since it incurs less profit loss than for the baseline case.

For the supply side, Figure 3.4 shows that the number of full-time drivers increases with B , while the number of part-time drivers decreases since they switch to full-time. Once the number of employees reaches the cap, the number of both full- and part-time drivers remain the same. From Figure 3.5, we observe that drivers work longer on average: full-time drivers work 26.5 hours per week, and part-time drivers work 9 hours. Since limited benefits are provided, the platform's capacity decreases. Thus, the platform needs to charge a relatively high price to match the demand and supply, which leads to a reduction in consumer surplus. With a higher price, part-time drivers' hourly wage increases and so does their welfare. Full-time drivers' welfare decreases because of the cap, and the overall drivers' welfare is lower than the baseline case. In the Figure 3.5 welfare comparison, the driver classification was not consistent for the two cases as not all potential switchers became full-time in the second case. Figure 3.6 defines driver categories that are consistent between cases, i.e., those who work full-time regardless of the benefits, potential switchers who only work full-time, and those who always work part-time. The welfare of drivers who always work full-time decreases with B , but under the new policy, the welfare slightly decreases first and then increases. Moreover, their welfare is higher than the baseline when the benefit amount is high

enough. The welfare of the potential switchers increases with B in both cases. And the welfare under the new policy is lower than the baseline since not all switchers receive the benefits. Lastly, the new policy does not affect the welfare of drivers who always work part-time: the welfare is the same in both cases and decreases with B .

For lower benefit amounts, social welfare for the new policy is also lower than the baseline. However, as the benefits increase, the difference between the two policies becomes smaller. And as the benefits are high enough, the new policy generates higher social welfare.

Figure 3.7 and Figure 3.8 present how the profit and the portion fulfilled by employees change with \bar{n} when the platform provides different benefits packages ($B=150$ and $B=200$). From Figure 3.7, we can observe that the platform's profit is a non-increasing function of \bar{n} . Thus, without any regulation, the platform will choose a sufficiently low threshold to optimize its profit. Then the question becomes how to determine a proper threshold for the policymakers. As \bar{n} decreases, fewer orders are fulfilled by employees (see Figure 3.8). For example, as $\bar{n} = 150$, over 90% of the orders are fulfilled by employees, while as $\bar{n} = 70$, only half of the orders are fulfilled by employees. Potentially, the policymaker can set a minimum proportion of the workforce to consist of full-time employees to avoid the platform choosing a lower \bar{n} . However, this approach should be viewed as a starting point. It is possible policymakers can derive other policy mechanisms to achieve the same effect. It is important, however, to give the platforms a certain degree of control over the makeup of their workforce in order to achieve a reasonable profit margin.

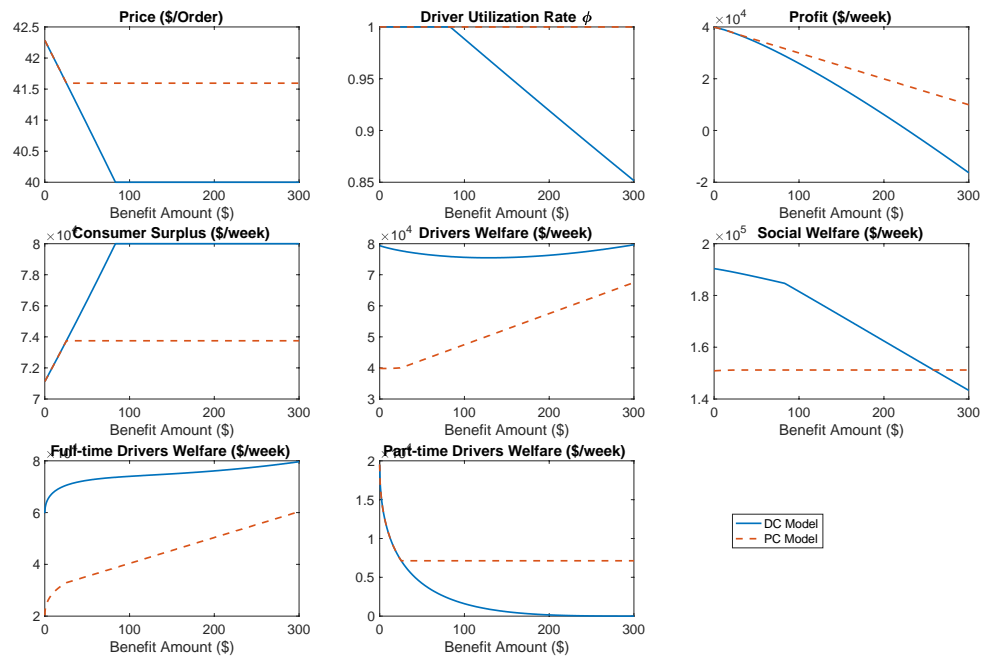


Figure 3.4 DC vs. PC Model

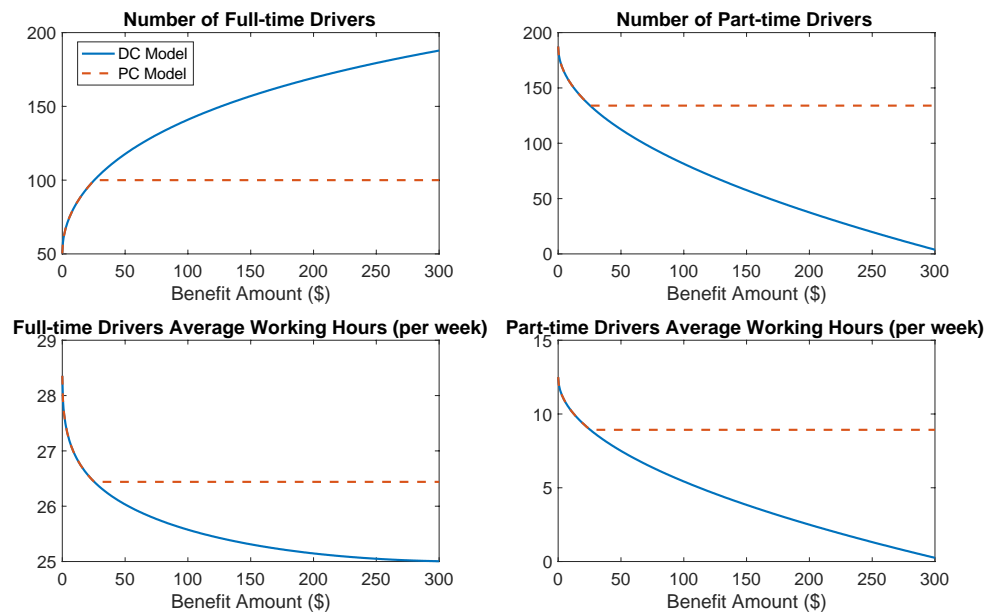


Figure 3.5 Supply: DC vs. PC Model

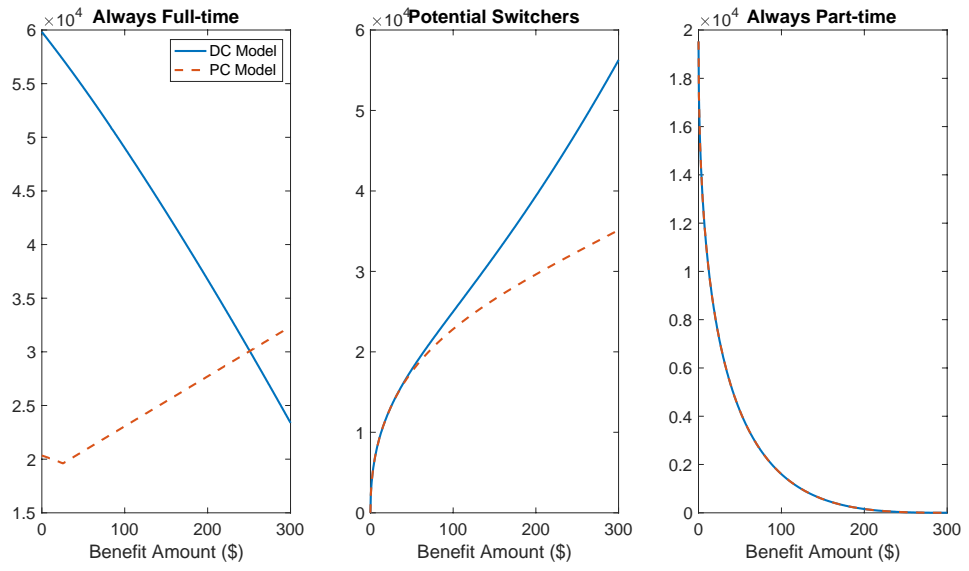


Figure 3.6 Drivers Welfare: DC vs. PC Model

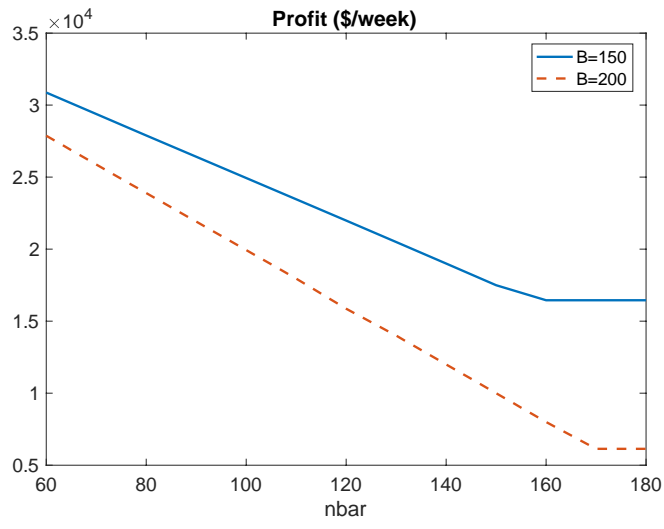


Figure 3.7 Profit with Different \bar{n}

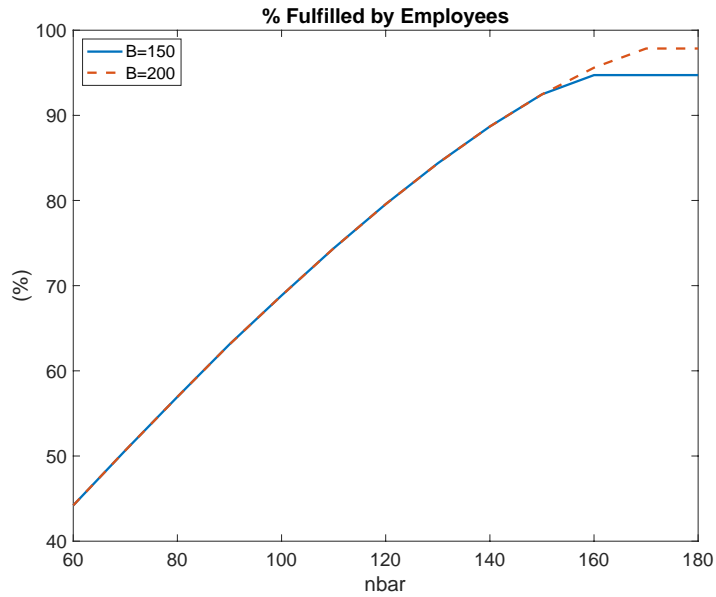


Figure 3.8 Percentage Fulfilled by Employees with Different \bar{n}

3.4 Conclusions and Future Research

In this chapter, we discuss a ride-hailing platform that adopts a hybrid workforce: the platform hires both employees and contractors and provides benefits to the employees. Currently, ride-hailing platforms, like Uber and Lyft, hire independent contractors only. Contractors have the flexibility to decide when to work and platforms have indirect control over them. Employees usually have a fixed schedule which guarantees a fixed capacity for platforms. By adopting a hybrid model, a platform can achieve some of the dynamic workforce characteristics it needs and also control profit loss (relative to an all-contractor workforce).

In our analytical model, the platform is a profit-maximizer and determines the optimal price. Drivers decide whether to work for the platform and how long to work. Under our baseline model, all drivers are independent contractors. The pricing strategy

for the platform is to charge either the unconstrained optimal price or the equilibrium price whichever maximizes its profit. Then we consider two models under the hybrid system. In the Driver-Controlled model, drivers who work over a certain number of hours are classified as employees and paid a benefits package. However, in the Platform-Controlled model, the platform can decide who to hire as employees and the number of employees is capped. In both models, the platform adopts a similar pricing strategy as the baseline case: choosing between the equilibrium price and the unconstrained optimal price. In these cases, the equilibrium price becomes a function of the benefit amount. With the benefits, more drivers are willing to work longer. We then have two types of employees: drivers who work full-time regardless of the benefits and those who only work full-time for the benefits (switchers). In the DC model, all switchers will be classified as employees, but only part of them are allowed to be employees in the PC model.

To better understand the influence of different policies and employment modes, we conduct several numerical studies. We first compare the baseline model and the DC model. We find that the platform maximizes its profit when there is no benefit. The profit decreases with the benefit amount. This is consistent with the strong gig company opposition to policy proposals that classify independent contractors as employees. As B increases, more drivers are willing to work as employees. Thus the platform's capacity grows and the driver utilization rate ϕ starts to be less than one. With smaller ϕ , the service quality of the platform increases since customers incur less waiting time on average. However, when ϕ is relatively low, it is possible that no driver is willing to work because of the low hourly earnings. So a lower bound for the utilization rate is

needed, for example, $\phi \geq 0.6$. In our current scenario, ϕ is between 0.85 and 1, which locates it in a reasonable range. In the DC model, employees are better off and their welfare is higher than with no benefits. However, part-time drivers receive a lower hourly wage, which lowers their welfare compared to the baseline model. Social welfare also decreases with the benefit amount.

The platform incurs a substantial profit loss under the DC model, for example, the profit drops 60% when B is 150. To balance the platform's profit and the drivers' welfare, we propose an alternative policy (PC model) that allows platforms to hire at most \bar{n} employees while maintaining a minimum proportion of orders fulfilled by the employees. In the PC model, the platform charges a higher price than the DC model because of the lower capacity. However, the profit increases over 30% when B is 150. Part-time drivers are also better off since their hourly earnings are higher than the DC model.

In this chapter, we consider a one-period model with fixed demand. In reality, the demand for ride-hailing platforms varies by season and the time of the day. For example, the demand is higher during rush hours and the holiday season. In this way, under a multi-period model, which includes different demand patterns, the platform may adopt different prices and staffing decisions accordingly. Such a model would demonstrate the benefits of retaining dynamism in the workforce, which maintaining a certain level of contractor-drivers could provide. Another potential model extension would be to consider more complex benefit packages. For example, there could be different benefit levels depending on the hours worked. Also, our model include a fixed benefit amount for full-time employees – a portion of the benefits could increase with

hours worked / earnings level. In our model, drivers' opportunity cost is a quadratic function of working hours, i.e., $ch + \frac{1}{2}\theta h^2$. Here we assume c is uniformly distributed and θ is a constant. Potentially, the degree of drivers' aversion to working longer (θ) may vary among drivers. A more complicated model with variation in θ could be analyzed to capture the diversity of drivers better. We hope future research can address these and other issues.

Appendices

APPENDIX A: Chapter 2

A.1 Wait Time and Labor Size

In this section, we investigate the relationship between wait time and the number of drivers using the following regression model:

$$\log(WT_{kmt}) = \beta \log(N_{kmt}) + \gamma X_t + \Phi_m + T_k + \varepsilon_{kmt} \quad (\text{A.1})$$

where WT_{kmt} represents the average wait time of market m at time t for tier k ; N_{kmt} is the number of in-network drivers in tier k at market m and at time t ; X_t represents time-related dummies, including weekdays, rush hours and a month dummy; Φ_m and T_k are the market and time fixed effects; and ε_{kmt} is the customer-based random effect with mean zero. Table A.1 shows the regression results of Equation A.1. After controlling for the market and time fixed effects, we observe that more in-network drivers lead to less wait time. Thus, retaining more drivers on the platform will help decrease the wait time and will further affect customer utility. Orders for higher tiers will also take less wait time.

Table A.1 Regression Results on Wait Time

	EST (SE)
Log(NumDriver)	-0.149*** (0.012)
Premium	-0.010 (0.012)
Luxury	-0.363*** (0.019)
Market Fixed Effect	Yes
Time Fixed Effect	Yes
Observations	249,929

*** p<0.001; ** p<0.01; * p<0.05; + p<0.1

A.2 ROI of the Bonus Policy

In this section, we choose different investment amounts for the bonuses provided to the retained in-network drivers in Section 2.6.1. We increase the bonus budget from 40% to 80% of the marketing cost of price discounts and calculate the ROI of bonuses for different investments. Table A.2 shows that the ROI of bonuses is around 1.6 to 1.7, and always lower than the ROI of discounts and upgrades.

Table 12 The ROI of the Bonuses

	EST (SE)
40%	1.682 (0.240)
50%	1.656 (0.239)
60%	1.662 (0.232)
70%	1.645 (0.230)
80%	1.638 (0.223)

A.3 Derivation of the FOCs

From Equation 2.4, we know that in the first stage, the reservation utility for consumer i is:

$$\begin{aligned}
 u_{imt} &= \alpha_1 C_{imt} + \alpha_2 \log(\text{Dist}_{imt}) + E[\max\{U_{ikmt}\}] + \Phi_{mt}^u + \varepsilon_{imt}^u \\
 &= \alpha_1 C_{imt} + E[\max\{U_{ikmt}\}] + X_{imt} + \varepsilon_{imt}^u
 \end{aligned}
 \tag{A.2}$$

where X_{imt} includes trip characteristics such as the trip distance, weekday dummy as well as market-time fixed effects.

Combining Equation A.2 and 2.5, we can calculate the probability that consumer i make a reservation, which is denoted as p_{imt} :

$$\begin{aligned}
p_{imt} &= \frac{\exp(\alpha_1 C_{imt} + E[\max\{U_{ikmt}\}] + X_{imt})}{1 + \exp(\alpha_1 C_{imt} + E[\max\{U_{ikmt}\}] + X_{imt})} \\
&= \frac{\exp(\alpha_1 C_{imt} + X_{imt}) \exp[\log(\sum_{k=1}^3 e^{\delta_{ikmt}})]}{1 + \exp(\alpha_1 C_{imt} + X_{imt}) \exp[\log(\sum_{k=1}^3 e^{\delta_{ikmt}})]} \\
&= \frac{\exp(\alpha_1 C_{imt} + X_{imt}) \sum_{k=1}^3 e^{\delta_{ikmt}}}{1 + \exp(\alpha_1 C_{imt} + X_{imt}) \sum_{k=1}^3 e^{\delta_{ikmt}}}
\end{aligned} \tag{A.3}$$

Here we focus on one market m at a specific time t , so in the following context, we will omit the subscripts for market m and time t for simplification. From Equation A.3, we know that the probability of consumer i making a reservation is p_i . Also, for each consumer i , the probability of booking tier k given a reservation is $\frac{e^{\delta_{ik}}}{\sum_{j=1}^3 e^{\delta_{ij}}}$. We denote the total number of consumers at market m and time t as M . Then by aggregating the demand for all consumers, the market share for each tier k can be specified as follows:

$$d_k = \sum_{i=1}^M p_i * \frac{e^{\delta_{ik}}}{\sum_{j=1}^3 e^{\delta_{ij}}} = \sum_{i=1}^M \frac{e^{\alpha_1 C_{it} + X_i + \delta_{ik}}}{1 + e^{\alpha_1 C_{it} + X_i} \sum_{j=1}^3 e^{\delta_{ij}}} \tag{A.4}$$

Recall that the platform's coupon decision variable c is denoted as the average coupon value provided to all customers. For simplification and calculation tractably, we assume that as a customer making a stage-1 decision, the discount in her utility function equals c at any time t and for any customer i . Thus, we replace C_{it} with c in the following derivations.

In this way, the derivatives of each market share can be written as:

$$\begin{aligned}
\frac{\partial d_k}{\partial c} &= \sum_{i=1}^M \frac{\alpha_1 e^{\alpha_1 C_{it} + X_i + \delta_{ik}}}{(1 + e^{\alpha_1 C_{it} + X_i} \sum_{j=1}^3 e^{\delta_{ij}})^2} \\
\frac{\partial d_k}{\partial \theta_k} &= 0
\end{aligned}$$

(A.5)

On the supply side, we can omit mt and rewrite Equation 2.11 as follows:

$$\begin{aligned} N_1 &= K_1[(1 - \theta_1)d_1r_1]^{\eta_1} \\ N_2 &= K_2\{[\theta_1d_1 + (1 - \theta_2)d_2]r_2\}^{\eta_1} \\ N_3 &= K_3[(\theta_2d_2 + d_3)r_3]^{\eta_1} \end{aligned} \quad (\text{A.6})$$

where $K_j = \exp(\eta_0 + \eta_2 \log(N_k^T) + \eta_3 \text{tier} + \Phi^v + \xi_k)$ and $r_k = (1 - \gamma^{\text{in}})P_k \text{dist}$. Thus, we can obtain the following derivatives:

$$\begin{aligned} \frac{\partial N_1}{\partial \theta_1} &= -K_1\eta_1d_1r_1[(1 - \theta_1)d_1r_1]^{\eta_1-1} \\ \frac{\partial N_1}{\partial \theta_2} &= 0 \\ \frac{\partial N_1}{\partial c} &= K_1\eta_1(1 - \theta_1)r_1 \frac{\partial d_1}{\partial c} [(1 - \theta_1)d_1r_1]^{\eta_1-1} \\ \frac{\partial N_2}{\partial \theta_1} &= K_2\eta_1d_1r_2\{[\theta_1d_1 + (1 - \theta_2)d_2]r_2\}^{\eta_1-1} \\ \frac{\partial N_2}{\partial \theta_2} &= -K_2\eta_1d_2r_2\{[\theta_1d_1 + (1 - \theta_2)d_2]r_2\}^{\eta_1-1} \\ \frac{\partial N_2}{\partial c} &= K_2\eta_1r_2\left[\theta_1 \frac{\partial d_1}{\partial c} + (1 - \theta_2) \frac{\partial d_2}{\partial c}\right] \{[\theta_1d_1 + (1 - \theta_2)d_2]r_2\}^{\eta_1-1} \\ \frac{\partial N_3}{\partial \theta_1} &= 0 \\ \frac{\partial N_3}{\partial \theta_2} &= K_3\eta_1d_2r_3[(\theta_2d_2 + d_3)r_3]^{\eta_1-1} \\ \frac{\partial N_3}{\partial c} &= K_3\eta_1r_3\left(\theta_2 \frac{\partial d_2}{\partial c} + \frac{\partial d_3}{\partial c}\right) [(\theta_2d_2 + d_3)r_3]^{\eta_1-1} \end{aligned} \quad (\text{A.7})$$

By plugging Equation A.6 and A.7 into Equation 2.15, we can solve the equation system and derive the value of R_{kmt} .

Appendix B: Chapter 3

Proof of Lemma 1:

Assume that the demand of the platform equals to $a - kp$. We consider two cases: (1) demand is greater than or equal to supply, i.e., $a - kp \geq S$; (2) demand is less than or equal to supply, i.e., $a - kp \leq S$. For case 1, we are demand limited and thus have $\phi = 1$ so that platform profit is $\Pi = (1 - \delta) \cdot p \cdot S$. Since S is increasing in p , we can observe that Π is an increasing function of p . Since $a - kp$ decreases with p the highest profit is obtained when $a - kp = S$. In this way, we show that $a - kp \leq S$ always holds.

Proof of Proposition 1:

From Problem 1, we have $\Pi = (1 - \delta) \cdot p \cdot (a - kp) = (1 - \delta)(-kp^2 + ap)$, which is a concave function of p . So the optimal price for the unconstrained profit maximization problem is $p^c = \frac{a}{2k}$. The constraint $a - kp \leq \frac{N(\phi\delta p)^2}{2\theta\bar{c}}$ is equivalent to

$$\frac{N(\phi\delta)^2}{2\theta\bar{c}}p^2 + kp - a \geq 0, \quad \text{i.e.,} \quad \frac{-\sqrt{(k\theta\bar{c})^2 + 2N\theta\delta^2 a\bar{c}} - k\theta\bar{c}}{N\delta^2} \leq p \leq \frac{\sqrt{(k\theta\bar{c})^2 + 2N\theta\delta^2 a\bar{c}} - k\theta\bar{c}}{N\delta^2}.$$

Since $p \geq 0$, then the constraint is equivalent to $p \leq \frac{\sqrt{(k\theta\bar{c})^2 + 2N\theta\delta^2 a\bar{c}} - k\theta\bar{c}}{N\delta^2} := p^m$.

Thus, the optimal price equals $\max\{p^c, p^m\}$.

Proof of Proposition 2:

The optimal working hours for a driver with opportunity cost equals $h^* := h(c) = \frac{\phi\delta p - c}{\theta}$ when there are no benefits provided. Drivers with $h^* \geq h_0$, i.e., $c \in (0, \phi\delta p - \theta h_0]$, will work as employees regardless of the benefits. However, with the extra benefits, some drivers are willing to switch to full-time. We define $\Delta = h_0 - h^*$.

Then the driver's incremental utility for working hours h_0 , i.e., an additional Δ hours is $U(h^* + \Delta) - U(h^*)$, which can be calculated as follows:

$$\begin{aligned}
U(h^* + \Delta) - U(h^*) &= \phi\delta p(h^* + \Delta) - c(h^* + \Delta) - \frac{1}{2}\theta(h^* + \Delta)^2 + B - \phi\delta p h^* + c h^* + \frac{1}{2}\theta h^{*2} \\
&= \phi\delta p \Delta - c\Delta - \frac{1}{2}\theta\Delta^2 - \theta h^* \Delta + B \\
&= (\phi\delta p - c)\Delta - \theta\Delta \frac{\phi\delta p - c}{\theta} - \frac{1}{2}\theta\Delta^2 + B \\
&= B - \frac{1}{2}\theta\Delta^2
\end{aligned}$$

Drivers will work h_0 hours only if $U(h^* + \Delta) \geq U(h^*)$, i.e., $\Delta \leq \sqrt{\frac{2B}{\theta}}$. Also, $\Delta = h_0 - h^* = h_0 - \frac{\phi\delta p - c}{\theta}$. Then we have those with $\phi\delta p - \theta h_0 \leq c \leq \phi\delta p - \theta h_0 + \sqrt{2B\theta}$ will work as employees if benefits are provided. Drivers with $\phi\delta p - \theta h_0 + \sqrt{2B\theta} \leq c \leq \phi\delta p$ will work as contractors since $U(h_0) \leq U(h^*)$.

Derivation of S for DC Model:

The platform's capacity is:

$$\begin{aligned}
S &= NE[h^* \mathbb{1}\{c \in (0, \phi\delta p - \theta h_0)\}] + NE[h_0 \mathbb{1}\{c \in (\phi\delta p - \theta h_0, \phi\delta p - \theta h_0 + \sqrt{2B\theta})\}] \\
&\quad + NE[h^* \mathbb{1}\{c \in (\phi\delta p - \theta h_0 + \sqrt{2B\theta}, \phi\delta p)\}] \\
&= N \int_0^{\phi\delta p - \theta h_0} \frac{\phi\delta p - c}{\theta} \frac{1}{\bar{c}} dc + N \int_{\phi\delta p - \theta h_0}^{\phi\delta p - \theta h_0 + \sqrt{2B\theta}} h_0 \frac{1}{\bar{c}} dc + N \int_{\phi\delta p - \theta h_0 + \sqrt{2B\theta}}^{\phi\delta p} \frac{\phi\delta p - c}{\theta} \frac{1}{\bar{c}} dc \\
&= \frac{N(\phi\delta p)^2 - N(\theta h_0)^2}{2\theta\bar{c}} + \frac{N h_0 \sqrt{2B\theta}}{\bar{c}} + \frac{N(\theta h_0 - \sqrt{2B\theta})^2}{2\theta\bar{c}} \\
&= \frac{N(\phi\delta p)^2}{2\theta\bar{c}} + \frac{NB}{\bar{c}}
\end{aligned}$$

Proof of Proposition 3:

Noting that the first constraint is equivalent to $\phi \leq 1$, we can write the Lagrangian equation as:

$$L(p, \phi) = (1 - \delta)p(a - kp) - \frac{NB(\phi\delta p - \theta h_0 + \sqrt{2B\theta})}{\bar{c}} + u(1 - \phi) + v\left(\phi - \frac{a - kp}{S}\right) \quad (2.3)$$

Then we can list the KKT conditions:

$$\frac{\partial L(p, \phi)}{\partial p} = (1 - \delta)(a - 2kp) - \frac{NB\phi\delta}{\bar{c}} + \frac{v}{S}\left(k + \phi \frac{\partial S}{\partial p}\right) = 0$$

$$\frac{\partial L(p, \phi)}{\partial \phi} = -\frac{NB\delta p}{\bar{c}} - u + v\left(1 + \frac{\phi}{S} \frac{\partial S}{\partial \phi}\right) = 0$$

$$u(1 - \phi) = 0$$

$$a - kp \leq S$$

$$\phi = (a - kp)/S$$

$$u \geq 0$$

Then we discuss two subcases:

(i) When $\phi = 1$, $a - kp = \frac{N(\phi\delta p)^2}{2\theta\bar{c}} + \frac{NB}{\bar{c}}$. In the manner similar to the proof of

proposition 2, we can derive the price as $p = \frac{\sqrt{(k\theta\bar{c})^2 + 2N\theta\delta^2(a\bar{c} - NB)} - k\theta\bar{c}}{N\delta^2} = p^m$

(ii) When $\phi < 1$, $u = 0$. Then we have:

$$(1 - \delta)(a - 2kp) - \frac{NB\phi\delta}{\bar{c}} + \frac{v}{S}\left(k + \phi \frac{\partial S}{\partial p}\right) = 0$$

$$-\frac{NB\delta p}{\bar{c}} + v\left(1 + \frac{\phi}{S} \frac{\partial S}{\partial \phi}\right) = 0$$

$$\phi = (a - kp)/S$$

From the first two equations, we have

$$\left[\frac{NB\phi\delta}{\bar{c}} - (1 - \delta)(a - 2kp)\right]\left(S + \phi \frac{\partial S}{\partial \phi}\right) = \frac{NB\delta p}{\bar{c}}\left(k + \phi \frac{\partial S}{\partial p}\right)$$

$$\frac{NB\phi\delta}{\bar{c}}\left(S + \phi \frac{\partial S}{\partial \phi}\right) - \frac{NB\delta p}{\bar{c}}\left(k + \phi \frac{\partial S}{\partial p}\right) = (1 - \delta)(a - 2kp)\left(S + \phi \frac{\partial S}{\partial \phi}\right)$$

$$\frac{NB\phi\delta}{\bar{c}}\phi \frac{\partial S}{\partial \phi} - \frac{NB\phi\delta}{\bar{c}}p \frac{\partial S}{\partial p} + \frac{NB\phi\delta}{\bar{c}}S - \frac{NB\delta p}{\bar{c}}k = (1 - \delta)(a - 2kp)\left(S + \phi \frac{\partial S}{\partial \phi}\right)$$

$$\frac{NB\phi\delta}{\bar{c}}S - \frac{NB\delta p}{\bar{c}}k = (1-\delta)(a-2kp)\left(S + \phi \frac{\partial S}{\partial \phi}\right)$$

$$\frac{NB\delta}{\bar{c}}(\phi S - pk) = (1-\delta)(a-2kp)\left(S + \phi \frac{\partial S}{\partial \phi}\right)$$

$$\frac{NB\delta}{\bar{c}}(a-2kp) = (1-\delta)(a-2kp)\left(S + \phi \frac{\partial S}{\partial \phi}\right)$$

$$(a-2kp)\left[(1-\delta)\left(S + \phi \frac{\partial S}{\partial \phi}\right) - \frac{NB\delta}{\bar{c}}\right] = 0$$

which is equivalent to:

$$(a-2kp)\left[(1-\delta)\frac{3N(\phi\delta p)^2}{2\theta\bar{c}} + \frac{NB(1-2\delta)}{\bar{c}}\right] = 0$$

We first show that $(1-\delta)\frac{3N(\phi\delta p)^2}{2\theta\bar{c}} + \frac{NB(1-2\delta)}{\bar{c}} > 0$. As $\delta \leq 0.5$, $(1-\delta)\frac{3N(\phi\delta p)^2}{2\theta\bar{c}} + \frac{NB(1-2\delta)}{\bar{c}} \geq (1-\delta)\frac{3N(\phi\delta p)^2}{2\theta\bar{c}} > 0$. When $\delta > 0.5$, $(1-\delta)\frac{3N(\phi\delta p)^2}{2\theta\bar{c}} + \frac{NB(1-2\delta)}{\bar{c}} > 0$ is

equivalent to:

$$(1-\delta)\frac{3N(\phi\delta p)^2}{2\theta\bar{c}} > \frac{NB(2\delta-1)}{\bar{c}}$$

$$3(1-\delta)(\phi\delta p)^2 > (2\delta-1)2B\theta$$

$$\sqrt{\frac{3-3\delta}{2\delta-1}} \phi\delta p > \sqrt{2B\theta}$$

In this chapter, we are focusing the case when the platform has positive number of full-time and part-time drivers, which requires $\phi\delta p - \theta h_0 > 0$ and $\phi\delta p - \theta h_0 + \sqrt{2B\theta} < \phi\delta p$, i.e., $\phi\delta p > \theta h_0$ and $\theta h_0 > \sqrt{2B\theta}$. In this way, we have $\phi\delta p >$

$\sqrt{2B\theta}$. Meanwhile, $\sqrt{\frac{3-3\delta}{2\delta-1}}$ is a decreasing function of δ . As $0.5 < \delta \leq 0.8$, we have

$\sqrt{\frac{3-3\delta}{2\delta-1}} \in [1, +\infty)$, i.e., $\sqrt{\frac{3-3\delta}{2\delta-1}} \geq 1$. Thus, $\sqrt{\frac{3-3\delta}{2\delta-1}} \phi\delta p > \sqrt{2B\theta}$ holds. Thus $p =$

$$\frac{a}{2k} := p^c.$$

Proof of Proposition 4:

From the proof of Proposition 2, drivers with $c \in (0, \phi\delta p - \theta h_0]$ will work as employees regardless of the benefits, and those with $c \in (\phi\delta p - \theta h_0, \phi\delta p - \theta h_0 + \sqrt{2B\theta}]$ want to work full-time if benefits are provided. The platform chooses the most “enthusiastic” drivers as employees, i.e., those with the lowest opportunity cost c . Switchers’ opportunity costs are uniformly distributed between $\phi\delta p - \theta h_0$ and $\phi\delta p - \theta h_0 + \sqrt{2B\theta}$. And the fraction of switchers becoming employees is t . Thus, the total number of potential switchers with the lowest c are classified as employees, i.e., $c \in (\phi\delta p - \theta h_0, \phi\delta p - \theta h_0 + t\sqrt{2B\theta}]$.

Derivation of S for PC Model:

The platform’s capacity is:

$$\begin{aligned}
 S &= NE[h^*\mathbb{1}\{c \in (0, \phi\delta p - \theta h_0)\}] + NE[h_0\mathbb{1}\{c \in (\phi\delta p - \theta h_0, \phi\delta p - \theta h_0 + t\sqrt{2B\theta})\}] \\
 &\quad + NE[h^*\mathbb{1}\{c \in (\phi\delta p - \theta h_0 + t\sqrt{2B\theta}, \phi\delta p)\}] \\
 &= N \int_0^{\phi\delta p - \theta h_0} \frac{\phi\delta p - c}{\theta} \frac{1}{\bar{c}} dc + N \int_{\phi\delta p - \theta h_0}^{\phi\delta p - \theta h_0 + t\sqrt{2B\theta}} h_0 \frac{1}{\bar{c}} dc + N \int_{\phi\delta p - \theta h_0 + t\sqrt{2B\theta}}^{\phi\delta p} \frac{\phi\delta p - c}{\theta} \frac{1}{\bar{c}} dc \\
 &= \frac{N(\phi\delta p)^2 - N(\theta h_0)^2}{2\theta\bar{c}} + \frac{N(\theta h_0 - t\sqrt{2B\theta})^2}{2\theta\bar{c}} + \frac{Nt h_0 \sqrt{2B\theta}}{\bar{c}} \\
 &= \frac{N(\phi\delta p)^2}{2\theta\bar{c}} + \frac{NBt^2}{\bar{c}}
 \end{aligned}$$

Driver Welfare Calculations

DC Model Welfare:

We denote the full-time drivers’ welfare as DW_f and the part-time drivers’ welfare as DW_p , then we have:

$$DW_f = \frac{N}{\bar{c}} \int_0^{\phi\delta p - \theta h_0} \left[\frac{(\phi\delta p - c)^2}{2\theta} + B \right] dc + \frac{N}{\bar{c}} \int_{\phi\delta p - \theta h_0}^{\phi\delta p - \theta h_0 + \sqrt{2B\theta}} \left[\phi\delta p h_0 - c h_0 - \frac{1}{2} \theta h_0^2 + B \right] dc$$

$$\begin{aligned}
&= \frac{N}{\bar{c}} \left[\frac{1}{6\theta} c^3 - \frac{\phi\delta p}{2\theta} c^2 + \frac{(\phi\delta p)^2}{2\theta} c + Bc \right]_{\phi\delta p - \theta h_0}^{\phi\delta p - \theta h_0} + \frac{N}{\bar{c}} \left[\left(\phi\delta p h_0 - \frac{1}{2} \theta h_0^2 + B \right) c - \frac{h_0}{2} c^2 \right]_{\phi\delta p - \theta h_0}^{\phi\delta p - \theta h_0 + \sqrt{2B\theta}} \\
&= \frac{N}{\bar{c}} \left[\frac{(\phi\delta p - \theta h_0)^3}{6\theta} - \frac{\phi\delta p}{2\theta} (\phi\delta p - \theta h_0)^2 + \frac{(\phi\delta p)^2}{2\theta} (\phi\delta p - \theta h_0) + B(\phi\delta p - \theta h_0) \right] \\
&\quad + \frac{N}{\bar{c}} \left[\left(\phi\delta p h_0 - \frac{1}{2} \theta h_0^2 + B \right) \sqrt{2B\theta} - \frac{h_0}{2} (2\phi\delta p - 2\theta h_0 + \sqrt{2B\theta}) \sqrt{2B\theta} \right] \\
DW_p &= \frac{N}{\bar{c}} \int_{\phi\delta p - \theta h_0 + \sqrt{2B\theta}}^{\phi\delta p} \left[\frac{(\phi\delta p - c)^2}{2\theta} \right] dc = \frac{N}{\bar{c}} \left[\frac{1}{6\theta} c^3 - \frac{\phi\delta p}{2\theta} c^2 + \frac{(\phi\delta p)^2}{2\theta} c \right]_{\phi\delta p - \theta h_0 + \sqrt{2B\theta}}^{\phi\delta p} \\
&= \frac{N}{\bar{c}} \left[\frac{(\phi\delta p)^3}{6\theta} - \frac{(\phi\delta p - \theta h_0 + \sqrt{2B\theta})^3}{6\theta} + \frac{\phi\delta p}{2\theta} (\phi\delta p - \theta h_0 + \sqrt{2B\theta})^2 - \frac{(\phi\delta p)^2}{2\theta} (\phi\delta p - \theta h_0 + \sqrt{2B\theta}) \right]
\end{aligned}$$

Then $DW = DW_f + DW_p$.

PC Model Welfare:

Different from the DC Model, only switchers with $c \in (\phi\delta p - \theta h_0, \phi\delta p - \theta h_0 + t\sqrt{2B\theta}]$ are classified as employees. Thus, driver welfare becomes:

$$\begin{aligned}
DW_f &= \frac{N}{\bar{c}} \int_0^{\phi\delta p - \theta h_0} \left[\frac{(\phi\delta p - c)^2}{2\theta} + B \right] dc + \frac{N}{\bar{c}} \int_{\phi\delta p - \theta h_0}^{\phi\delta p - \theta h_0 + t\sqrt{2B\theta}} \left[\phi\delta p h_0 - c h_0 - \frac{1}{2} \theta h_0^2 + B \right] dc \\
&= \frac{N}{\bar{c}} \left[\frac{1}{6\theta} c^3 - \frac{\phi\delta p}{2\theta} c^2 + \frac{(\phi\delta p)^2}{2\theta} c + Bc \right]_0^{\phi\delta p - \theta h_0} + \frac{N}{\bar{c}} \left[\left(\phi\delta p h_0 - \frac{1}{2} \theta h_0^2 + B \right) c - \frac{h_0}{2} c^2 \right]_{\phi\delta p - \theta h_0}^{\phi\delta p - \theta h_0 + t\sqrt{2B\theta}} \\
&= \frac{N}{\bar{c}} \left[\frac{(\phi\delta p - \theta h_0)^3}{6\theta} - \frac{\phi\delta p}{2\theta} (\phi\delta p - \theta h_0)^2 + \frac{(\phi\delta p)^2}{2\theta} (\phi\delta p - \theta h_0) + B(\phi\delta p - \theta h_0) \right] \\
&\quad + \frac{N}{\bar{c}} \left[\left(\phi\delta p h_0 - \frac{1}{2} \theta h_0^2 + B \right) t\sqrt{2B\theta} - \frac{h_0}{2} (2\phi\delta p - 2\theta h_0 + t\sqrt{2B\theta}) t\sqrt{2B\theta} \right] \\
DW_p &= \frac{N}{\bar{c}} \int_{\phi\delta p - \theta h_0 + t\sqrt{2B\theta}}^{\phi\delta p} \left[\frac{(\phi\delta p - c)^2}{2\theta} \right] dc = \frac{N}{\bar{c}} \left[\frac{1}{6\theta} c^3 - \frac{\phi\delta p}{2\theta} c^2 + \frac{(\phi\delta p)^2}{2\theta} c \right]_{\phi\delta p - \theta h_0 + t\sqrt{2B\theta}}^{\phi\delta p} \\
&= \frac{N}{\bar{c}} \left[\frac{(\phi\delta p)^3}{6\theta} - \frac{(\phi\delta p - \theta h_0 + t\sqrt{2B\theta})^3}{6\theta} + \frac{\phi\delta p}{2\theta} (\phi\delta p - \theta h_0 + t\sqrt{2B\theta})^2 \right. \\
&\quad \left. - \frac{(\phi\delta p)^2}{2\theta} (\phi\delta p - \theta h_0 + t\sqrt{2B\theta}) \right]
\end{aligned}$$

Bibliography

Ailawadi, Kusum L. and Sunil Gupta (2014), *Sales promotions*, in Russel S. Winter and Scott A. Neslin (Eds), *The History of Marketing Science*. Hackensack, NJ: World Scientific Publishing.

Albuquerque, Paulo, Polykarpos Pavlidis, Udi Chatow, Kay-Yut Chen, and Zainab Jamal (2012), “Evaluating Promotional Activities in an Online Two-Sided Market of User-Generated Content,” *Marketing Science*, 31(3), 406–432.

Armstrong, Mark (2006), “Competition in Two-Sided Markets,” *The RAND Journal of Economics*, 37(3), 668–691.

Bai, Jiaru, Kut C. So, Christopher S. Tang, Xiqun (Michael) Chen, and Hai Wang (2018), “Coordinating Supply and Demand on an On-Demand Service Platform with Impatient Customers,” *Manufacturing & Service Operations Management*, 21(3):556-570.

Banerjee Siddhartha, Ramesh Johari, and Carlos Riquelme (2015), “Pricing in ride-sharing platforms: A queueing-theoretic approach.” *Proceedings of the sixteenth ACM Conference on Economics and Computation*, pages 639

Baron, David P. (2018) “Disruptive Entrepreneurship and Dual Purpose Strategies: The Case of Uber,” *Strategy Science*, 3(2):439–462

Bawa, Kapil and Robert Shoemaker (2004), “The Effects of Free Sample Promotions on Incremental Brand Sales,” *Marketing Science*, 23(3), 345–363.

Bawa, Kapil, Srini S. Srinivasan, and Rajendra K. Srivastava (1997), “Coupon Attractiveness and Coupon Proneness: A Framework for Modeling Coupon Redemption,” *Journal of Marketing Research*, 34(4), 517–525.

Benjaafar, Saif, Jian-Ya Ding, Guangwen Kong, and Terry Taylor (2021), “Labor Welfare in On-Demand Service Platforms,” *Manufacturing & Service Operations Management*, 24(1):110-124.

Benjaafar, Saif, and Ming Hu (2019), “Operations Management in the Age of the Sharing Economy: What Is Old and What Is New?” *Manufacturing & Service Operations Management*, 22(1):93-101.

Bimpikis, Kostas, Ozan Candogan, and Daniela Saban (2019), “Spatial Pricing in Ride-Sharing Networks,” *Operations Research*, 67(3):744-769.

Cachon, Gérard P., Kaitlin M. Daniels, and Ruben Lobel (2017), “The Role of Surge Pricing on a Service Platform with Self-Scheduling Capacity,” *Manufacturing & Service Operations Management*, 19(3):368-384.

Cahuc, Pierre, Stéphane Carcillo, and André Zylberberg (2014), *Labor Economics*, MIT press

Caillaud, Bernard and Bruno Jullien (2003), “Chicken & Egg: Competition among Intermediation Service Providers,” *The RAND Journal of Economics*, 34(2), 309–328.

Castillo, Juan Camilo, Dan Knoepfle, and Glen Weyl (2017), “Surge pricing solves the wild goose chase,” *Proceedings of the 2017 ACM Conference on Economics and Computation*, ACM, pp. 241–242.

Chakravarty, Amiya K., (2021), “Blending Capacity on a Rideshare Platform: Independent and Dedicated Drivers,” *Production and Operations Management*, *Production and Operations Management Society*, vol. 30(8), pages 2522-2546

- Chen, Yiwei, and Ming Hu (2019), “Pricing and Matching with Forward-Looking Buyers and Sellers,” *Manufacturing & Service Operations Management*, 22(4):717–734.
- Chiara, Farronato, Jessica Fong, and Andrey Fradkin (2020), “Dog Eat Dog: Measuring Network Effects Using a Digital Platform,” *NBER Working Paper Series*.
- Chu, Junhong and Puneet Manchanda (2016), “Quantifying Cross and Direct Network Effects in Online Consumer-to-Consumer Platforms,” *Marketing Science*, 35(6), 870–893.
- Clements, Matthew T. and Hiroshi Ohashi (1994), “Indirect Network Effects and the Product Cycle: Video Games in the U.S., 1994-2002,” *The Journal of Industrial Economics*, 53(4), 515–542.
- Dong, Jing, and Rouba Ibrahim (2020), “Managing Supply in the On-Demand Economy: Flexible Workers, Full-Time Employees, or Both?” *Operations Research*, 68(4):1238-1264.
- Dubé, Jean-Pierre, Günter J. Hitsch, and Pradeep K. Chintagunta (2010), “Tipping and Concentration in Markets with Indirect Network Effects,” *Marketing Science*, 29(2), 216–249.
- Feliz Leon, Luis. (2022). “Here's What's in the New Bill Jointly Backed by Uber and the Teamsters in Washington State,” *Labor Notes*.
- Gurvich, Itai, Lariviere Martin, and Moreno Antonio (2017), “Operations in the On-Demand Economy: Staffing Services with Self-Scheduling Capacity,” *Sharing Economy: Making Supply Meet Demand*, Ming Hu (ed), *Springer Series in Supply Chain Management*, pp 249–278

Hagiu, Andrei, and Julian Wright (2019), “The status of workers and platforms in the sharing economy,” *Journal of Economics and Management Strategy*, 28(1): 97–108

Hall, Jonathan V., and Alan B. Krueger (2018), “An Analysis of the Labor Market for Uber’s Driver-Partners in the United States,” *ILR Review*, 71(3):705–732

Heiman, Amir, Bruce McWilliams, Zhihua Shen, and David Zilberman (2001), “Learning and Forgetting: Modeling Optimal Product Sampling Over Time,” *Management Science*, 47(4), 532–546.

Hu, Ming and Yun Zhou (2020) “Price, Wage, and Fixed Commission in On-Demand Matching,” Available at SSRN: <https://ssrn.com/abstract=2949513>

Kabra, Ashish, Elena Belavina, and Karan Girotra (2017), “The Efficacy of Incentives in Scaling Marketplaces,” *Working Paper*.

Kannan, P.K., and Gordon P. Wright (1991), “Modeling and Testing Structured Markets: A Nested Logit Approach,” *Marketing Science*, 10(1), 58-82

Li, Hongshuang, Sanjay Jain, and P.K. Kannan (2019), “Optimal Design of Free Samples for Digital Products and Services,” *Journal of Marketing Research*, 56(3), 419-438

Lobel, Ilan, Sebastien Martin, and Haotian Song (2021), “Employees, Contractors, or Hybrid: An Operational Perspective,” *NYU Stern School of Business Forthcoming*, Available at SSRN: <https://ssrn.com/abstract=3878215>

Ming, Liu, Tunay I. Tunca, Yi Xu, and Weiming Zhu (2019), “An Empirical Analysis of Market Formation, Pricing, and Revenue Sharing in Ride-Hailing Services,” *Working Paper*.

- Neslin, Scott A. (1990), "A Market Response Model for Coupon Promotions," *Marketing Science*, 9(2), 125–125.
- Parker, Geoffrey G., Marshall W. Van Alstyne, and Sangeet Paul Choudary (2016), *Platform Revolution*. New York, NY: W.W. Norton & Company.
- Reibstein, David J. and Phillis A. Traver (1982), "Factors Affecting Coupon Redemption Rates," *Journal of Marketing*, 46(4), 102–113.
- Rochet, Jean-Charles and Jean Tirole (2003), "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1, 990–1029.
- Sudhir, K. and Nathan Yang (2018), "The Value of Free Upgrades," *Working Paper*.
- Ward, Ronald W. and James E. Davis (1978), "Coupon Redemption," *Journal of Advertising Research*, 18(4), 51–58.
- Taylor, Terry A. (2018), "On-Demand Service Platforms," *Manufacturing & Service Operations Management*, 20(4):704-720.
- Yu, Jiayi Joey, Christopher S. Tang, Zuo-Jun Max Shen, and Xiquan Michael Chen (2019), "A Balancing Act of Regulating On-Demand Ride Services," *Management Science*, 66(7):2975-2992.
- Zhang, Dennis J., Hengcheng Dai, Lingxiu Dong, Fangfang Qi, Nannan Zhang, Xiaofei Liu, Zhongyi Liu, Jiang Yang (2020), "The Long-term and Spillover Effects of Price Promotions on Retailing Platforms: Evidence from a Large Randomized Experiment on Alibaba," *Management Science*, 66(6), 2589–2609.
- Zheng, Jinyang, Fei Ren, Yong Tan, and Xi Chen (2020). "Optimizing Two-Sided Promotion for Transportation Network Companies: A Structural Model with Conditional Bayesian Learning," *Information Systems Research*, 31(3), 692–714.