

MANAGING SELF-SCHEDULING CAPACITY

Kaitlin M. Daniels

A DISSERTATION

in

Operations, Information and Decisions

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2017

Supervisor of Dissertation

Gerard P. Cachon, Professor of Operations, Information and Decisions

Graduate Group Chairperson

Catherine Schrand, Celia Z. Moh Professor, Professor of Accounting

Dissertation Committee

Noah Gans, Professor of Operations, Information and Decisions

Ruben Lobel, Data Scientist, Airbnb

Xuanming Su, Professor of Operations, Information and Decisions

Senthil Veeraraghavan, Professor of Operations, Information and Decisions

MANAGING SELF-SCHEDULING CAPACITY

© COPYRIGHT

2017

Kaitlin Marie Daniels

This work is licensed under the
Creative Commons Attribution
NonCommercial-ShareAlike 3.0
License

To view a copy of this license, visit

<http://creativecommons.org/licenses/by-nc-sa/3.0/>

ACKNOWLEDGEMENT

I would like to thank: the Wharton Risk Management Center and the Mack Institute for Innovation Management for their generous financial support; my committee members for their patience and constructive feedback; Ian Eaves for his true partnership in this endeavor; Richard Daniels for his expertise and his willingness to read; Kathleen Engelmann for her constant encouragement and her restrained prodding; Alex Daniels for his perspective on struggle; Alix Barasch, Nora Becker, Sunita Desai, Ana Gazmuri, Jess Jeffers, Cinthia Konichi, Emma Levine, Kathy Li, Ellie Prager, and Preethi Rao for being a source of balance, inspiration, and joy.

ABSTRACT

MANAGING SELF-SCHEDULING CAPACITY

Kaitlin M. Daniels

Gerard P. Cachon

Gig-economy platform like Uber, Lyft, Postmates, and Instacart have created markets in which independent service providers provide on-demand service to consumers. A hallmark of this arrangement is that providers decide for themselves when, where, and how much to work. In other words, the platform does not set its capacity's schedule; instead its capacity "self-schedules." This decentralization of decision making can create value for providers. The platform's challenge is then to devise a contract with its capacity that allows it to capture some of this value. I study the platform's contracting problem in three chapters. In the first, I show that the platform can benefit from allowing its providers to self-schedule. In the second, I study the platform's strategy when coordinating supply and demand across multiple states of the world. I show that the resulting dynamic pricing policy can be beneficial to consumers, despite widespread dislike of the real-world practice. I also show that, in many cases, the platform need not independently vary payments to providers to achieve near-optimal profit. Instead the platform may pay its providers a fixed percent commission on the price paid by consumers per completed service. In the final chapter, I argue that the findings above are distinct from the traditional two-sided markets literature. Though a classic two-sided market model experiences near-optimal performance of the fixed commission in many cases, the market conditions that produce poor fixed commission performance differ between the gig-economy model and the two-sided markets model. Because the two-sided market model does not accurately predict poor gig-economy fixed commission performance, it is important to study a model tailored the gig-economy to understand gig-economy specific applications.

TABLE OF CONTENTS

ACKNOWLEDGEMENT	iii
ABSTRACT	iv
LIST OF TABLES	vii
LIST OF ILLUSTRATIONS	viii
PREFACE	ix
CHAPTER 1 : Demand Response in Electricity Markets: Voluntary and Automated Curtailment Contracts	1
1.1 Introduction	1
1.2 Literature Review	5
1.3 Consumer Behavior	8
1.4 Firm's Contract Choice	10
1.5 System Welfare and Environmental Impact	14
1.6 Optimal Curtailment Contracts	16
1.7 Numerical Results	19
1.8 Conclusion	21
CHAPTER 2 : The Role of Surge Pricing on a Service Platform with Self-Scheduling Capacity	24
2.1 Introduction	24
2.2 Literature Review	28
2.3 Model	30
2.4 Contract Design	35
2.5 Fixed Participation Cost	41

2.6	Numerical Study	44
2.7	Self-Scheduling vs. Central-Scheduling of Capacity	52
2.8	Discussion	54
2.9	Conclusion	55
CHAPTER 3 : Distinguishing the Gig-Economy from Two-Sided Markets		58
3.1	Introduction	58
3.2	Model	62
3.3	Gig Economy	65
3.4	Two-Sided Markets	67
3.5	Numerical Analysis	70
3.6	Conclusion	74
APPENDIX		75
BIBLIOGRAPHY		130

LIST OF TABLES

TABLE 1 :	Linear Contract Parameters	11
TABLE 2 :	Tested parameter values. All combinations of these values constitute 14,700 scenarios.	45
TABLE 3 :	Frequency of different versions of the fixed contract.	46
TABLE 4 :	Relative Profitability of Suboptimal Contracts. Profit performance of the four suboptimal contracts relative to the optimal contract in all 14,700 scenarios (left table) and in the 2,253 scenarios in which the fixed contract serves both demand states (right table). The subscripts f , w , p , β , and o to refer to the fixed, dynamic wage, dynamic price, commission and optimal contracts, respectively. . .	47
TABLE 5 :	Relative Consumer Surplus with Poor Utilization. The ratio of con- sumer surplus and number of providers with the dynamic wage, dy- namic price, commission or optimal contract to the fixed contract in the 591 scenarios with poor utilization.	50
TABLE 6 :	Relative Consumer Surplus with Poor Service. The ratio of con- sumer surplus and number of providers with the dynamic wage, dy- namic price, commission or optimal contract to the fixed contract in the 1,662 scenarios with poor service.	51
TABLE 7 :	A summary of tested parameter values. All combinations of these values constitute 104,976 numerical experiments.	70
TABLE 8 :	Quartiles and mean of P_g , P_t and P_g/P_t	71

LIST OF ILLUSTRATIONS

FIGURE 1 : Structure of the Market with Curtailment Contracts 4

FIGURE 2 : This figure illustrates the firm’s and the social planner’s contract choices as functions of market price sensitivity and cost variance. In Region A $\Pi_A^L \geq \Pi_V^L$ and $W_A^L \geq W_V^L$, in Region B $\Pi_A^L \geq \Pi_V^L$ but $W_V^L \geq W_A^L$, and in Region C $\Pi_V^L \geq \Pi_A^L$ and $W_V^L \geq W_A^L$ 15

FIGURE 3 : A comparison of contract performance for the firm’s profit 20

FIGURE 4 : A comparison of contract performance for the social welfare. 21

FIGURE 5 : Timeline of events 30

FIGURE 6 : An example of demand and capacity rationing with a fixed contract. 36

FIGURE 7 : Relative Consumer Surplus as a Function of Commission. The ratio of consumer surplus with the commission contract to consumer surplus with the fixed contract as a function of the commission earned by providers with the commission contract in the 2,253 scenarios in which the fixed contract serves both demand states. Squares indicate the average ratio for scenarios grouped by the commission contract commission in 0.05 intervals. 52

FIGURE 8 : Illustrative plots. All plots have $\mu_g = 1$, $CV_g = .5$, and $CV_f = .25$. In the first row, curves are defined by $\{\gamma, \rho, \mu_f\}$, where red has $\{.1, .3, 5\}$; blue has $\{.1, .7, 5\}$; green has $\{.6, .3, 5\}$; black has $\{.1, .3, 1\}$. In the second row, curves are defined by $\{\alpha, \rho, \mu_f\}$, where red has $\{.5, .3, 5\}$; blue has $\{.5, .7, 5\}$; green has $\{.1, .3, 5\}$; black has $\{.5, .3, 1\}$. In the third row, curves are defined by $\{\alpha, \gamma, \mu_f\}$ where red has $\{.5, .1, 5\}$; blue has $\{.5, .5, 5\}$; green has $\{.1, .1, 5\}$; black has $\{.5, .1, 1\}$ 72

PREFACE

Here I study markets in which a firm relies on capacity to provide a good or service that it does not directly control. This capacity is the aggregation of individual producers or service providers who produce/serve according to their own self-interest. These individuals face non-trivial, non-constant, and private opportunity costs for the time spent producing/serving through the firm, so each individual only provides capacity to the firm if his earnings from doing so exceed his opportunity cost. In other words, the firm does not set its capacity's service/production schedule, instead its capacity "self-schedules."

Self-scheduling arrangements are most associated with the on-demand service industry commonly known as the gig-economy. Gig-economy firms serve as platforms that connect customers with self-scheduling service providers. Noteworthy examples include the ride-sharing platforms Uber and Lyft, the delivery platforms Postmates and Instacart, and the on-demand labor platforms TaskRabbit and Bellhops. According to Farrell and Greig (2016), between 2012 and 2015 the fraction of JPMorgan Chase account holders receiving monthly income from participation in the gig-economy grew 10-fold, constituting 1% of account holders in 2015. Furthermore, the proportion of account-holders having ever received income from participation in the gig-economy grew 47-fold, indicating that 4.2% of account holders worked a gig at some point during that time period (Farrell and Greig, 2016).

Providers report that their primary attraction to the gig-economy is their ability to self-schedule. A recent Benenson Strategy Group survey of Uber drivers reports that 87% of respondents partner with Uber "to be my own boss and set my own schedule" (Hall and Krueger, 2015a). The key consequence of this framework is that decentralization decision making can create value for providers. However, it does not immediately follow that this arrangement also benefits the platform or consumers.

In three chapters, I study the firm's contracting problem. First I pursue the most fundamental question: whether self-scheduling can be more profitable than a traditional capacity

arrangement. Second I derive a firm’s profit optimal dynamic pricing contract and determine its effect on consumer welfare. Finally, I illustrate the differences between the gig-economy and the traditional two-sided markets literature.

The first chapter is motivated by a setting with self-scheduling capacity that existed long before the gig-economy: the electricity market. “Curtailed contracts” in electricity markets allow a firm to pay electricity consumers *not* to use electricity and to sell this foregone consumption as “virtual” generation on the electricity market. In the firm’s contract with electricity consumers, it may either allow consumers to curtail their consumption as they wish (i.e. self-schedule) or it may require consumer to commit to a level of curtailment. In this context, the firm evaluates the profitability of self-scheduling in a market where the firm is a price taker and price decreases in the production of the firm. I find that the firm generally prefers self-scheduling when consumers face sufficiently varied opportunity costs and when the slope of the price curve is sufficiently small.

The second chapter is motivated by the dynamic pricing strategies used by some gig-economy firms (most notably Uber’s “Surge Pricing”). The firm chooses a pricing strategy to coordinate its self-scheduling capacity with variable demand. The firm’s profit maximizing incentive structure charges consumers a demand-contingent price and offers providers a demand-contingent wage. I consider the profit and consumer surplus implications of restricting prices to be constant across all demand regimes and, alternatively, allowing demand-contingent prices but restricting wages to be a fixed percentage of prices (i.e. providers earn a fixed commission). I find that while restricting the firm to a fixed price causes significant profit loss for the firm, requiring a fixed commission generally produces near-optimal profit. I further demonstrate that consumers can benefit from demand-contingent prices due to depressed prices in low demand regimes and expanded access to service in high demand regimes.

In the final chapter, I argue that the findings described above are distinct from predictions made by the traditional two-sided markets literature. In that literature, membership on the

platform happens on a much longer time scale than in the gig-economy, so the classic models ignore capacity constraints that are very relevant to the operations of the gig-economy. For example, a retailer chooses to accept American Express based on the volume of customers that carry American Express. In contrast, an Uber driver can only serve one customer at a time, so even if there are many customers demanding rides at the same time, the Uber driver only cares that he is assured a passenger. I find that these difference lead the two models to make significantly different predictions. In particular, I focus on the profitability of a fixed commission in each setting. I show that though a classic two-sided market model experiences near-optimal performance of the fixed commission in many cases, the market conditions that produce poor fixed commission performance differ between the gig-economy model and the two-sided markets model. Because the two-sided market model does not accurately predict poor gig-economy fixed commission performance, it is important to study a model tailored the gig-economy to understand gig-economy specific applications.

The goal of this work is to guide practitioners, regulators, and theoreticians in their understanding of self-scheduling capacity. I demonstrate that practitioners can profit from a self-scheduling arrangement, and that the challenge of providing sufficient capacity when facing variable demand and heterogeneous workers can be overcome through dynamic prices and wages. I additionally show regulators that the profitability of self-scheduling does not preclude benefit to consumers and workers. For example, consumers can benefit from dynamic prices precisely because workers self-schedule - dynamic prices can improve consumers' access to service at busy times while decreasing the time workers spend idle during slow times. Furthermore, I demonstrate to theoreticians the importance of considering the details that distinguish the gig-economy from classic examples of two-sided markets.

CHAPTER 1 : Demand Response in Electricity Markets: Voluntary and Automated Curtailment Contracts

1.1. Introduction

Electricity markets today suffer from a fundamental flaw: consumption does not respond to market signals. This is a result of the organization of the electricity supply chain, as illustrated in Figure 1. Conventional generators submit a menu of prices and production levels to the market's clearinghouse, the Independent System Operator (ISO). The ISO constructs the market supply curve by ordering these bids from least to most expensive per kilowatt hour (kWh). Market demand determines the price at which all consumption is purchased, which is known as a uniform-price auction. However, consumers do not satisfy their demand by purchasing electricity on the market themselves. Instead, consumers are served by local utilities and other electricity retailers that buy on the market at the market price and sell to consumers at a fixed fee per kWh. Though the market price varies, end-consumers only experience their fixed fee, even in times of peak load. As a result, consumers that would otherwise be priced out of the market continue to consume, and demand can creep dangerously close to the system's limit, threatening black- and brown-outs.

There are two well studied ways of addressing the insulation of consumption from market prices. The first, called real time pricing, removes the middle man and allows consumers to pay the market price. Alternatively, curtailment contracts, the topic of interest of this paper, insert a new middle man called the Curtailment Service Provider (CSP), who pays consumers not to consume and sells the induced foregone consumption on the market as *virtual supply* (see Figure 1). The CSP acts as a generator in the electricity market, only the electricity it sells is not newly generated but is instead electricity made available by the curtailment of contracted consumers' demand. Curtailment contracts have been shown to achieve the same efficiency as real time pricing (Chao and Wilson, 1987) and, as a result of a recent policy change, are growing in popularity as a demand management tool. In

2011 the Federal Energy Regulatory Commission mandated that foregone end-consumption be treated as conventional generation during peak load events, meaning the extra supply created by curtailed consumption can be sold on the market for the full market price.¹ In the next 5 years curtailment contracts, along with other demand response efforts, are projected to decrease peak demand by more than 4%.¹ In the next 5 years curtailment contracts, along with other demand response efforts, are projected to decrease peak demand by more than 4%.²

When a consumer curtails his consumption, he incurs an opportunity cost which varies based on the value of his initial consumption. For example, a residential consumer's value of air conditioning is increasing in the ambient temperature, so his opportunity cost from curtailing his electricity is higher during a heat wave than on a normal summer day. Under a curtailment contract, the CSP calls upon contracted consumers to reduce their consumption when market conditions make selling virtual load a profitable endeavor. Typically, the CSP participates the electricity market during peak load events, caused by extreme demand or generation outages. Because consumers do not anticipate peak load events, the value of the curtailed consumption is uncertain ex ante. Furthermore, a consumer's value of consumption is private information.

Curtailment contracts can be partitioned into two classes. Traditionally consumers relinquish control over their curtailment decisions, allowing the CSP to remotely adjust their consumption during peak load events. We will call this an automated curtailment contract. The curtailment amount automatically imposed by the CSP is determined by the consumer in advance of any particular peak load event. Therefore, curtailment under an automated contract is determined in expectation of the consumer's value of consumption at the time of curtailment. Technological advances in smart metering have introduced a new generation of curtailment contracts which we will call voluntary. Under the voluntary contract consumers decide the level of curtailment provided to the CSP at the time of the peak event.

¹FERC Order 745 - <http://www.ferc.gov/EventCalendar/Files/20110315105757-RM10-17-000.pdf>

²EIA 2011 report - <http://www.eia.gov/todayinenergy/detail.cfm?id=650>

The voluntary contract, therefore, allows consumers to make curtailment decisions with full knowledge of their cost.

In this paper, we study a CSP's choice of contract class to offer a set of consumers. In practice, CSPs offer distinct contracts to different customer segments (e.g. residential, commercial). A contract is designed for a particular subset of consumers who have similar attributes. We assume that the consumers within this subset have a homogeneous value distribution of electricity consumption. For example, interior heating and cooling represents nearly a third of the residential electricity consumption in the United States ³. The value of this consumption varies as a function of temperature, which affects consumers in the same geographic area equally.

The CSP's choice between automated and voluntary contracts echoes the trend in the service industry toward allowing service providers to self-schedule. Just as firms like Uber allow service providers to choose when and how much to work, a voluntary curtailment contract allows consumers to choose whether and how much to curtail their loads. Like electricity consumers curtailing their loads, self-scheduling service providers encounter an opportunity cost when they offer their services. The burgeoning operations literature on this topic considers providers with independent and identically distributed opportunity costs (e.g. Cachon et al. (2017), Gurvich et al. (2015), Ibrahim and Arifoglu (2015)). Curtailment contracts offer an appropriate application to extend this literature by studying a self-scheduling environment in which opportunity costs are correlated.

We characterize a CSP's choice of contract class as a function of market conditions and consumer characteristics. We find that the voluntary contract's value to the firm can be characterized by the difference between the expected curtailed load under the voluntary and automated contracts for a given payment; for a sufficiently large difference the firm prefers the voluntary contract. The sensitivity of market price to changes in consumption drives how large this difference must be; the larger the market sensitivity, the larger the difference must be for the firm to choose the voluntary contract. In the specific case of a symmetric

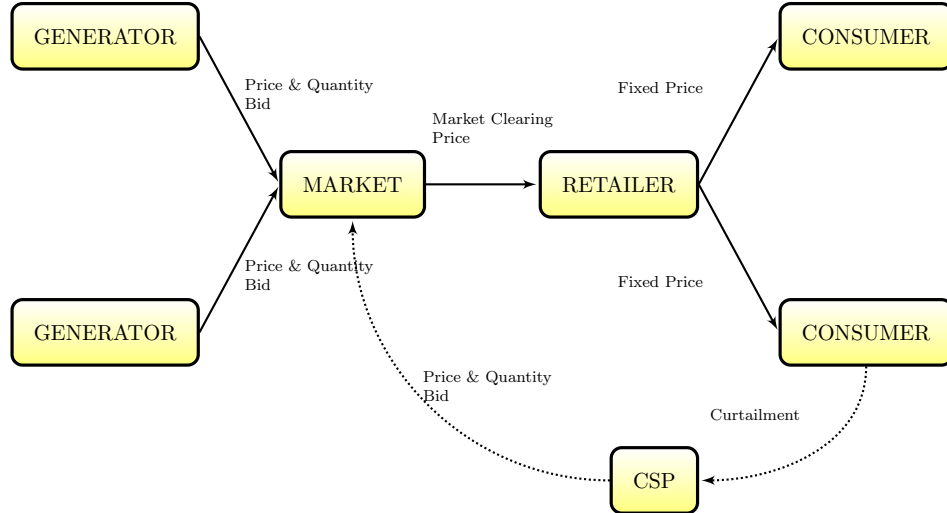


Figure 1: Structure of the Market with Curtailment Contracts

cost distribution, the firm’s contract decision is driven simply by the variance of cost. This application also allows us to measure the environmental effect of the CSP’s contract choice. While the profit maximizing CSP contract choice also maximizes the positive environmental impact of demand response, the incentives of the CSP and a welfare maximizing social planner are not aligned. We characterize the market and customer conditions that create this misalignment, those in which a social planner should encourage voluntary contracts even though the CSP would rather adopt the automated contract.

Below we frame the aforementioned analysis in terms of the relevant literature (Section 2). In Section 3, we discuss the consumer’s problem and in Section 4 the firm’s contract choice. In Section 5, we analyze the impact of the CSP’s contract choice on market coordination. Though many curtailment contracts observed in the market are linear in nature, we are interested in the firm’s optimal choice. Therefore, in Section 6 we relax the linear payment restriction and analyze the unrestricted “optimal” contract. To relate the optimal and linear contracts we study a numerical example in Section 7. We verify our results under the linear contract and find that the linear contract profit is a good approximation of the optimal contract profit. To streamline the presentation, we relegate all proofs to the Appendix.

1.2. Literature Review

The management of non-storable goods with periodic demand, such as electricity, has long been of interest to economists. The peak load pricing literature studies capacity and pricing decisions made by a firm facing demand that periodically shifts from low to high. For a thorough review please see Crew et al. (1995). When demand is stochastic, a firm with finite capacity faces a non-zero probability that its supply will be insufficient to satisfy demand during peak load events. In the electricity setting this means that a utility will interrupt its service to a fraction of its customers. To efficiently allocate its electricity, a utility may segment the service it provides based on the reliability of the service. Many papers have been dedicated to characterizing the optimal menu of price and reliability contracts to offer under this service differentiation based on reliability. See Chao and Wilson (1987), Marchand (1974), Chao et al. (1986), Smith (1989), Oren and Doucet (1990), and Oren and Smith (1992). These contracts are the direct ancestors of the automated curtailment contracts studied in this paper. Both sets of contracts allow utilities to reduce demand to address insufficient supply. However, the incentives facing the firm in our model differ from those listed above. Previous papers have considered contracts that maximize social welfare. In contrast, our model incorporates the monetary incentives for efficient demand-side management, so our firm is purely a profit maximizer. Furthermore, previous models have assumed service delivery to be binary; a consumer's electricity needs are either met or not. The ubiquity of smart meters today makes the question of how much electricity to curtail from a consumer both relevant and important. We incorporate this dimension by including a continuous load curtailment decision variable.

Iyer et al. (2003) study an electricity supplier's decision to incentivise demand postponement. The supplier faces a two period decision horizon and has the option to delay a proportion of demand until the second period. This postponement comes at an exogenous cost per unit of demand postponed. The authors study how postponement influences the supplier's capacity investment decisions in both periods. The structure of the supplier's

interaction with consumers resembles the curtailment contracts studied in our model. As in our model, Iyer et al. (2003) must remunerate consumers in return for their participation in the demand management program. However, unlike our model the authors do not allow the supplier to determine how best to reward consumer participation. Their reimbursement is linear and exogenous.

Our work is born of a rich literature studying demand-side management. Our key innovation is the introduction of the voluntary curtailment contracts and the characterization of when it should be preferred over the automated type. Both the reliability-based service differentiation literature and Iyer et al. (2003) assume that all demand management uses contracts that we consider “automated.” Specifically, service reliability is determined before service interruption and interruption is controlled by the firm. Similarly, the engineering literature on demand response has primarily focused on how to optimally activate a portfolio of automated load curtailment customers (see for example Goyal et al. 2013 and Taylor and Mathieu 2014). Wu and Kapuscinski (2013) show how curtailment of intermittent generation can also benefit the system by reducing uncertainty for conventional generation sources. While under reliability-based service contracts consumers may sort themselves efficiently given their ex-ante understanding of their value for electricity, we expect consumers to have a better understanding of their valuation at the time they wish to consume electricity. The automated structure affords consumers no flexibility at the time of consumption. Allowing consumers to adjust their service curtailment based on new information at the time of consumption certainly generates extra value for consumers. In the analysis that follows we will show when that extra value can also benefit the firm.

In parallel to the supply-chain literature, one can consider the CSP firm as a retailer, while the customer is the supplier of curtailment. Under uncertain supplier cost, the flexibility of a voluntary supply contract can benefit both players. Kim and Netessine (2013) show how to design contracts to encourage collaboration in the supply chain. There are several contracts designed to create supply flexibility for firms facing stochastic demand (backup agreements

as in Eppen and Iyer (1997) and quantity flexible contracts as in Tsay and Lovejoy (1999)). Barnes-Schuster et al. (2002) show that these contracts are special cases of a larger problem in which the firm can purchase options from a supplier at a specified exercise price. The resulting chain coordinating contract is a piece-wise linear exercise price. However, flexibility is not necessarily profit improving when suppliers are not forced to comply with capacity requests from a manufacturer who has private information about demand. Cachon and Lariviere (2001) show that a manufacturer inducing capacity investment in a supplier, based on asymmetric information about demand, must perform a costly act (increase capacity requested or committed order size, lump sum transfer) to induce the appropriate capacity investment in the supplier. Similarly, suppliers in our chain are not forced to comply with capacity requests under the voluntary contract. In the analysis that follows we will show when supply flexibility is beneficial to the firm and when the firm is better off enforcing compliance.

The cost uncertainty and the flexibility of the voluntary contract also introduces uncertainty into the supply of the firm. The firm designs the voluntary contract to extract curtailment from consumers to sell on the market. Because the firm decides the contract structure before the consumer decides his curtailment level, the firm bears the risk of the uncertain curtailment supply. Issues of supply uncertainty have been thoroughly explored in the operations literature, see Yano and Lee (1995) for a review. Uncertainty mitigation strategies generally fall into two categories: inventory strategies and sourcing strategies. Inventory strategies hedge against stock-outs by increasing inventory levels. For example, Kim et al. (2010) show that a performance based contract can encourage the supplier to increase inventories and mitigate supply risk. Sourcing strategies balance supplier procurement costs with reliability (more reliable suppliers are more expensive). Since electricity is not practically storable, inventory strategies are infeasible in our setting. Instead ours is fundamentally a question of sourcing. How should a firm decide between offering an automated contract (reliable source) and a voluntary contract (unreliable source)? Yang et al. (2012), Babich et al. (2007), and Wu and Babich (2012) delineate scenarios in which a firm is willing to sac-

rifice reliability to achieve lower procurement cost or to allocate risk appropriately. Tomlin (2006) studies how the characteristics of the unreliable source, in terms of % uptime and expected disruption length, determine a firm’s choice between a reliable and an unreliable source with fixed procurement costs. Yang et al. (2009) investigates a firm’s willingness to pay for reliability information when a supplier’s reliability is private information. Similarly, Wang et al. (2010), Kim (2011), and Liu et al. (2010) explore a firm’s decision to invest in endeavors that will improve the reliability of its supplier. This parallels the premium we show that the firm pays to consumers under an automated contract.

1.3. Consumer Behavior

In this section we describe the behavior of a customer who has decided to participate in a contract offered by the CSP, henceforth referred to as the firm. Suppose the firm has an exogenously determined number of consumers, N , who are enrolled in the firm’s curtailment contract. Each consumer will be called upon by the firm during peak load events to reduce his energy consumption in return for payment. We would like to emphasize that the dynamics of our model are the opposite of a conventional supply chain; here the firm buys from the consumer to sell on the market instead of the other way around. In most applications the firm contracts with a pool of consumers. We characterize the time until the next event as a single period. We assume that consumers have a baseline value of curtailment, q , defined by

$$V(q) = -Cq^2$$

where the parameter C governs the consumer’s value of the forgone consumption. Because consumers have a positive valuation of their baseline consumption, their value for curtailment is negative and increasing in the amount of consumption they forgo. The cost of a given level of curtailment, q , varies depending on the task that is interrupted by curtailment. To capture this, we allow C to vary over time. Furthermore, we assume that consumers are similar in the pattern of their valuation of consumption over time. For example, commercial consumers all value their ability to light their stores more during the holiday shopping sea-

son than during an off season, and residential consumers all increasingly value their ability to turn on their air conditioners the hotter the weather is. Hence we assume that consumers are homogeneous in the value coefficient, C , and that this value coefficient varies over time. Under a curtailment contract, the firm offers each consumer incentives to alter his consumption pattern. These incentives are offered in response to peak load events, which we assume consumers do not predict. Furthermore, the firm does not observe C , so from both firm and consumer point of view, this quantity is uncertain. We therefore define C to be a random variable, with density g , distribution G and finite support $\{C_L, C_H\}$. In our model both the firm and consumers have full knowledge of the distribution of C . Whether the consumers decide their actual curtailment with full knowledge of C distinguishes the two contracts studied in this paper.

1.3.1. Linear Automated Contracts

Under the automated contract the firm offers a linear payment, w , per unit of the consumer's load curtailment. The

firm also offers a fixed transfer per transaction, f , which may be positive or negative. When $f = 0$ the firm simply offers a wholesale price contract. A consumer's corresponding value function is

$$V_A(q) = E[wq - Cq^2 + f\mathbb{1}\{q > 0\}].$$

Given knowledge of (w, f) a consumer chooses to participate in the contract and must specify the load reduction quantity, q_A that he will provide to the

firm at the next peak load event. This decision is made in advance of the peak load event, so neither consumers nor the firm know the realization of C . In practice, firms offering automated contracts remotely control load reduction so, once the curtailment level is set, the consumer cannot change it. The consumer determines his consumption, q_A , to maximize

his value given the payment structure:

$$q_A = \operatorname{argmax} V_A(q) = \frac{w}{2E[C]}.$$

1.3.2. Linear Voluntary Contracts

In contrast with the automated contract, under the voluntary contract each consumer learns the realization of C *before* choosing their curtailment. A consumer's corresponding value function is

$$V_V(q) = wq - Cq^2 + f\mathbb{1}\{q > 0\}.$$

In response to the firm's offered payment, (w, f) , each consumer chooses his curtailment to maximize his value:

$$q_v = \operatorname{argmax} V_V(q) = \frac{1}{2C}.$$

1.4. Firm's Contract Choice

Given the consumer behavior described above, the

firm wishes to maximize her expected profit. Under a given contract type, the lever at her disposal is the payment level (w, f) . By reducing either w or f the firm lowers her cost. However, the

firm is restricted to payment schemes under which consumers will participate. Furthermore, we expect each consumer's load reduction choice to be an increasing function of the payment they receive. Lowering w also reduces the load reduction delivered to the firm by each consumer. This leaves the firm with a smaller amount of virtual generation to sell on the market at the market price, $P(q)$. Recall that the market price is an increasing function of aggregate demand, meaning the market price is a decreasing function of q , the demand removed from the market. The firm's optimal contract must balance these countervailing forces. We assume that the change in market price from the CSP's participation in the

market will not be large enough to demonstrate higher order effects, we approximate the market price curve as $P(q) = a - bq$. Let q_j^i denote consumer i 's curtailment under contract type j . We can characterize the firm's expected profit under contract type j as

$$\begin{aligned} \max_{w_j, f_j} \Pi(w_j, f_j) &= E \left[P \left(\sum_{i=1}^N q_j^i \right) \sum_{i=1}^N q_j^i - \sum_{i=1}^N w_j q_j^i - f_j \right] \\ &s.t. V_j(q_j^i) \geq 0 \forall i. \end{aligned}$$

Note that $q_j^i(w_j, f_j)$ is independent of f_j . As a result the firm uses f_A to extract all of the expected consumer surplus under the automated contract. However, because the consumer decides to participate after observing his cost under the voluntary contract, the voluntary contract only extracts all consumer surplus if the consumer's cost achieves its highest value. The optimal payments and profits are summarized in Table 1.

Table 1: Linear Contract Parameters

		Voluntary	Automated
Full Linear Model	$E[q^L(w_j^L)]$	$\frac{w_V^L}{2} E \left[\frac{1}{C} \right]$	$\frac{w_A^L}{2} \frac{1}{E[C]}$
	w_j^L	$\frac{aE[\frac{1}{C}]}{2E[\frac{1}{C}] + bN(\text{Var}(\frac{1}{C}) + E[\frac{1}{C}]^2) - 1/C_H}$	$\frac{aE[C]}{E[C] + bN}$
	$f_{j,2}^L$	$-\frac{(w_V^L)^2}{4C_H}$	$-\frac{(w_A^L)^2}{4E[C]}$
	Π^L	$\frac{a^2 E[\frac{1}{C}]^2 N}{4(2E[\frac{1}{C}] + bN \text{Var}(\frac{1}{C}) + bNE[\frac{1}{C}]^2 - 1/C_H)}$	$\frac{a^2 N}{4(E[C] + bN)}$
Wholesale Model	$E[q^W(w)]$	$\frac{w_V^W}{2} E \left[\frac{1}{C} \right]$	$\frac{w_A^W}{2} \frac{1}{E[C]}$
	w_j^W	$\frac{aE[\frac{1}{C}]}{2E[\frac{1}{C}] + bN(\text{Var}(\frac{1}{C}) + E[\frac{1}{C}]^2)}$	$\frac{aE[C]}{2E[C] + bN}$
	f_j^W	0	0
	Π^W	$\frac{a^2 E[\frac{1}{C}]^2 N}{4(2E[\frac{1}{C}] + bN(\text{Var}(\frac{1}{C}) + E[\frac{1}{C}]^2))}$	$\frac{a^2}{4(2E[C] + bN)}$

In Table 1 we use the superscript L to denote quantities under the optimal contract where the fixed payment is non-zero. This is to differentiate from the special case where the fixed payment is zero, which we refer to as the wholesale contract and denote with the superscript W . Notice that in the first row the expected value of the individual load reduction for a

given w_j and f_j under the voluntary contract is at least as large as the load reduction under the automated. This follows from Jensen's inequality which requires

$$E\left[\frac{1}{C}\right] \geq \frac{1}{E[C]}.$$

By Jensen's inequality we see that the optimal unit payment w_j^L is no larger under the voluntary than under the automated. These two rows confirm that the voluntary contract has option value to consumers. Consumers are willing to accept smaller payments for the same load provided to the firm in exchange for the flexibility of the voluntary contract. In other words, the firm must pay a premium for a certain (reliable) supply of curtailment. We include in Table 1 the special case wholesale contract where $f_j^W = 0$ as a reference for future analysis. While many contracts currently in use follow this format, it is easy to see that $\Pi_j^L \geq \Pi_j^W$ for either the voluntary or the automated. We will focus our analysis on the two-part linear payment scheme (denoted by the superscript L), returning to use the wholesale contract as a benchmark in the numerical experiments of Section 7. Of primary interest is the fourth row in Table 1. This row describes the expected optimal profit under the voluntary and automated contracts respectively. The difference between voluntary and automated contract profit is

$$\Pi_V - \Pi_A = E[C] - \frac{2}{E[1/C]} + \frac{1}{C_H E[1/C]} - bNc_v^2 \quad (1.1)$$

where cv is the coefficient of variation of $1/C$. The relative profitability of the two contract types under consideration is both a function of market conditions and customer characteristics. In particular, the market affects a firm's choice of contract type through b , the slope of the market price curve and N , the number of contracted consumers. Customers influence profitability through the distribution of their cost coefficient, C . The slope of the market price curve, b , changes how cost uncertainty affects profitability. On average, the voluntary contract receives more load reduction from each consumer for the same level of per-unit payment. However, the variance in the load reduction quantity delivered under the

voluntary contract also causes variance in the total load curtailed, which affects the market price faced by the firm. We can conclude from (1.1) that in sufficiently sensitive markets the risk of low market price from the variance in load curtailed under the voluntary contract hurts expected revenue more than the gains from paying consumers less. Similarly, the more consumers the firm contracts, the larger the impact variance has on the load curtailed under the voluntary contract. In choosing between these contracts, the firm weighs this price risk against the lower wages per unit curtailment achieved under the voluntary contract.

To understand the role of the consumer's distribution of C on contract performance, we must disentangle three statistics: $E[C]$, $E[1/C]$, and $Var(1/C)$. Fixing the average cost $E[C]$, we show in Theorem 1 the choice of contract depends on the Jensen Gap, which is the slackness of the Jensen inequality, $E[1/C] - 1/E[C]$.

Theorem 1. *For a given $E[C]$ there exists $\bar{y} \geq \bar{z}$ such that*

1. *for all $E[\frac{1}{C}] \geq \bar{y}$, $\Pi_A^L \leq \Pi_V^L$*
2. *for all $E[\frac{1}{C}] \leq \bar{z}$, $\Pi_A^L \geq \Pi_V^L$.*

For a fixed $E[C]$ decreasing $E[1/C]$ decreases the Jensen Gap: a sufficiently small Jensen's Gap favors the automated contract and a sufficiently large Jensen's Gap favors the voluntary contract. Notice that the Jensen Gap is the difference in expected curtailment under the voluntary and automated contracts for a given unit payment. The larger this difference, the larger the payment required by the automated contract to obtain the same curtailment as the voluntary. At some point, the voluntary contract allows the firm to offer a small enough payment that the cost savings under the voluntary contract outweigh the price risk the firm must bear.

Theorem 2. *For symmetric distribution G with a given mean, increasing the variance of C increases the Jensen Gap.*

For symmetric distributions, the intuition guiding the firm's decision simplifies. The more varied the consumer cost of curtailment, the more consumers value flexibility. The firm

partially internalizes this value through its per transaction fee, f_V , causing firms to prefer voluntary contracts for sufficiently varied distributions of C .

1.5. System Welfare and Environmental Impact

In this section, we will consider the contract choice decision from the perspective of an integrated system. Instead of maximizing profit, the centralized decision maker's objective is to maximize system welfare, i.e. the sum of the consumer and firm surplus. Transfers between entities are no longer considered, so system welfare may be expressed as

$$W_j = E \left[\left(a - b \sum_{i=1}^N q_j^i \right) \sum_{i=1}^N q_j^i - C \sum_{i=1}^N (q_j^i)^2 \right].$$

Using Table 1 yields:

$$W_A = \frac{a^2 N}{4(E[C] + bN)}$$

$$W_V = \frac{q^2 E[1/c]^2 N}{2(2E[1/C] + bNE[1/C^2] - 1/C_H)} - \frac{a^2 E[1/C]^2 N (bNE[1/C^2] + E[1/C])}{4(2E[1/C] + bNE[1/C^2] - 1/C_H)}$$

Note that the firm extracts all consumer surplus under the automated contract so $W_A^L = \Pi_A^L$. We can show, however, that welfare under the voluntary contract exceeds the value of the firm's profit. This demonstrates that consumers benefit from the flexibility of the voluntary contract. From this fact we derive the following theorem.

Theorem 3. 1. If $\Pi_V^L \geq \Pi_A^L$, then $W_V \geq W_A$.

2. If $W_A \geq W_V$, then $\Pi_A^L \geq \Pi_V^L$.

Theorem 3 demonstrates that the firm's and the central planner's incentives can be aligned, but it also indicates that this is not always the case. Figure 2 demonstrates the alignment, or lack thereof, of firm and social planner incentives. In Region A, both the firm and the centralized decision maker respond to the high market price sensitivity and prefer the automated contract. In Region C, the firm and system planner similarly both prefer the

voluntary contract in response to low market price sensitivity and high cost uncertainty. However, in Region B the firm chooses an automated contract where the central planner prefers the voluntary contract. We conclude that a regulator may be inclined to encourage the voluntary contract in markets where the ratio of price sensitivity and cost uncertainty is not sufficient for the firm to adopt this contract.




Figure 2: This figure illustrates the firm's and the social planner's contract choices as functions of market price sensitivity and cost variance. In Region A $\Pi_A^L \geq \Pi_V^L$ and $W_A^L \geq W_V^L$, in Region B $\Pi_A^L \geq \Pi_V^L$ but $W_V^L \geq W_A^L$, and in Region C $\Pi_V^L \geq \Pi_A^L$ and $W_V^L \geq W_A^L$.

From an environmental perspective, consider a regulator whose objective is to minimize energy consumption, which is to maximize the total load reduction delivered by the curtailment contract. As shown in the following theorem, in this case the firm and the social planner prefer the same contract.

Theorem 4. $\Pi_V^L \geq \Pi_A^L$ if and only if $E[q_V^L] \geq E[q_A^L]$.

This means the more profitable contract is the one that induces more actual expected load reduction. This can also be shown for the wholesale contract. As a result a conservationist

social planner's incentives are aligned with those of a profit maximizing firm's.

1.6. Optimal Curtailment Contracts

We have thus far assumed that the firm is restricted to paying the consumer linearly. While this is a common contract structure, we are interested in the firm's optimal choice. In this section we will relax the linear payment restriction and study outcomes when the firm offers payment $w(q)$ as a general function of q .

1.6.1. The Optimal Automated Contract

As before, the consumer determines his curtailment quantity to maximize his value from participating in the contract:

$$q_A^* \in \operatorname{argmax}_{q_A} w(q_A) - E[C]q_A^2.$$

The firm uses her understanding of the consumer's response to the payment function $w(q)$ to maximize her profit:

$$\begin{aligned} \max_{w(\cdot)} \Pi_A &= (a - bq_A^*)q_A^* - w(q_A^*) \\ \text{s.t. } w(q_A^*) - E[C](q_A^*)^2 &\geq 0. \end{aligned}$$

Clearly the firm will choose $w(q_A^*)$ so that the participation constraint is tight; increasing $w(q_A^*)$ hurts profit. The tightness of the participation constraint allows $w(q_A^*)$ to be written as

$$w(q_A^*) = E[C](q_A^*)^2$$

which expresses the payment for the consumer's chosen curtailment quantity as a function of that quantity. Because of this direct relationship, we may rewrite the firm's program as

a maximization over induced curtailment quantities instead of over payments:

$$\max_{q_A^*} \Pi_A = (a - bNq_A^*)Nq_A^* - E[C]N(q_A^*)^2.$$

Solving the first order condition of this concave function, we find the curtailment quantity that the firm wishes to prompt from the consumer is $q_A^O = \frac{a}{2(bN+E[C])}$. The payment that produces this curtailment quantity is $w(q_A^O) = E[C]\frac{a}{2(bN+E[C])}^2$, which yields the optimal profit, $\Pi_A^O = \frac{a^2N}{4(Nb+E[C])}$. This is the same expected firm profit produced under a linear contract where payment is the sum of a per unit payment, $w_A/q = \frac{aE[C]}{E[C]+bN}$, plus a fixed payment, $f_A = -\frac{a^2E[C]}{4(E[C]+bN)^2}$, where in this case the fixed payment is a fee paid by the consumer. This is precisely the optimal payment scheme outlined in the previous section, so we conclude that the linear automated contract achieves optimality.

1.6.2. Optimal Voluntary Contracts

Under the voluntary contract, the firm again solves:

$$\begin{aligned} \max_{w_V} \Pi_V^O &= (a - bNq_V^*)Nq_V^* - w_V(q_V^*)N \\ \text{s.t. } w_V - C(q_V^*)^2 &\geq 0 \quad \forall C \end{aligned}$$

where q_V^* is the consumer's best response to the firm's chosen payment menu. However, unlike the automated contract the consumer's curtailment quantity varies in his realization of C . Though this information is hidden from the firm, its determination of the consumer's curtailment quantity allows the firm to infer the consumer's cost. So instead of paying consumers a per-unit rate for curtailment, the firm can offer a menu of payments that correspond to different levels of curtailment. Let us call the optimal payment menu (w_V^O, q_V^O) . By the revelation principle (Myerson, 1981), the solution to the program above is payoff-equivalent to the solution to the following program which restricts the solution to truth-

revealing mechanisms:

$$\begin{aligned} \max_{w_V(), q_V()} \quad & E_C[(a - bNq_V(C))Nq_V(C) - w_V(C)N] \\ \text{s.t.} \quad & w_V(C) - Cq_V^2(C) \geq w_V(\hat{C}) - Cq_V^2(\hat{C}) \quad \forall C, \hat{C} \quad (\text{IC}) \\ & w_V(C) - Cq_V^2(C) \geq 0 \quad \forall C. \quad (\text{PC}) \end{aligned}$$

Under the linear contract the firm understands how the consumer should choose q_V as a function of his cost, C : the consumer's curtailment choice is inversely proportional to his cost. The incentive compatibility (IC) constraint requires that the optimal menu preserve the firm's understanding of the functional relationship between a consumer's curtailment choice and his cost and hence can infer the consumer's cost from his chosen curtailment quantity. The firm performs this profit maximization with the additional participation constraint (PC) that requires all types receive a positive payoff. The resulting payment menu and induced curtailment quantities as functions of C are outlined in the following theorem.

Theorem 5. *If the distribution of $-C$ has a non-decreasing failure rate then the optimal contract menu is*

$$q_V^O(C) = \frac{ag(C)}{2((b+C)g(C) + G(C))}; w_V^O(C) = \int_C^{C_H} q_V^O(s)^2 ds + Cq_V^O(C)^2,$$

which yields optimal expected profit

$$\Pi_V^O = \frac{a^2N}{4} E \left[\frac{1}{bN + C + \frac{G(C)}{g(C)}} \right]. \quad (1.2)$$

Notice that expected firm profit under the optimal contract resembles expected profit under the linear contract; profit is inversely related to the slope of the market price, b , and, holding all else constant, is increasing in $E[1/C]$. However, beyond these observations the optimal contract is not analytically tractable. We will therefore rely on numerical analysis in the

following section to understand how it relates to the linear contracts studied in previous sections. In particular, we will demonstrate that the expected profit of the two-part tariff is a good approximation of the profit earned under the optimal contract. We conclude that it is sufficient to study the two-part tariff to understand the benefits and disadvantages of voluntary curtailment contracts.

1.7. Numerical Results

To demonstrate the theoretical intuition established above, we report the results of a numerical experiment. We calibrate the market conditions of our experiment using market price data from PJM Interconnection, the Regional Transmission Organization serving the Mid-Atlantic region of the US. For simplicity, we assume a curtailment event to last 1 hour. Market payment for demand response is assumed to follow a linear model $P(q) = a - bq$, where $a = 300\$/MW$ and $b = .03\$/MW^2$ are estimated using PJM's May 2013 supply curve data. Additionally, we use a case study by EnerNOC³, a CSP, to calibrate the range of consumer curtailment costs.

EnerNOC has an automated contract agreement with Four Seasons, a large produce wholesaler whose temperature-controlled warehouse requires high electricity input. Four Seasons reportedly earned \$11,000 for its participation in the contract over 25 curtailment events averaging 400kW of load reduction during the winter events and 1MW of load reduction during the summer events. The case study includes no information about fixed payments, so we assume this contract is an wholesale contract for our cost calibration. Using the expressions for the optimal curtailment quantity under an automated contract we estimate Four Season's $E[C]$ to be approximately $15000\$/MW^2$. Given the limited information about the distribution of C we assume C to be normally distributed with mean $E[C]$ and standard deviation σ , truncated over the support $[1 + \frac{1500000}{\sigma}, 30001 - \frac{1500000}{\sigma}]$. We vary σ to observe it's impact on profit for EnerNOC and the welfare of the system.

³Four Seasons Produce Turns to EnerNOC for Fresh Ideas in Reducing Energy Use - <http://www.enernoc.com/our-resources/case-studies/four-seasons-produce-turns-to-enernoc-for-fresh-ideas-in-reducing-energy-use>




Figure 3: A comparison of contract performance for the firm’s profit

Figure 3 reports the profit performance of each contract studied above for varied values of the standard deviation of C . From Theorem 2, we know that the variance of a symmetric distribution is directly related to the Jensen gap, which affects the profit difference between contract types. Figure 3 demonstrates this result: for sufficiently high variance the voluntary profit curve crosses the automated profit curve for both the full linear and wholesale models. Notice that the two-part tariff linear voluntary profit mirrors the optimal voluntary profit. At its best the two-part tariff achieves 99.9% of the optimal profit, and the two profits are at their closest in the region in which the firm would pick a voluntary contract over an automated contract. As in Cachon and Zhang (2006), we find that the linear contract is a practical substitute for the optimal voluntary contract.

From a social planner’s perspective, Figure 4 shows how a firm’s contract decision affects the system welfare. Without fixed payments, both the firm and the social planner prefer the voluntary contract; firm and social planner incentives are aligned. In contrast, using two-part linear payments, the social planner prefers the voluntary contract for a broader set of consumer types than the firm. Unlike under the wholesale contract, the full two-part tariff linear contract fails to coordinate firm and social planner incentives.

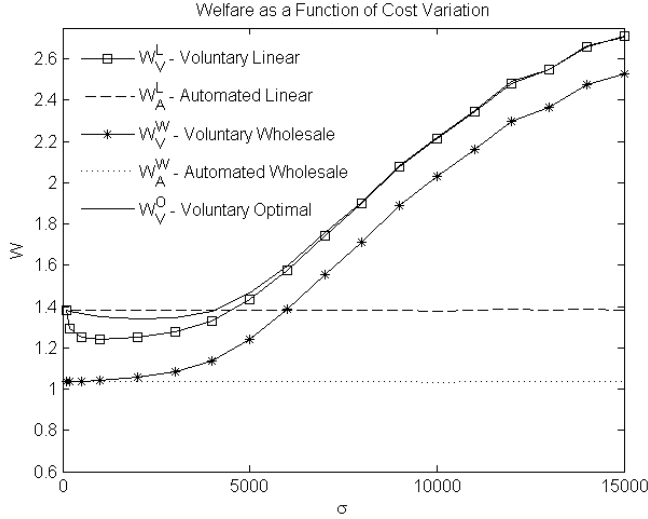


Figure 4: A comparison of contract performance for the social welfare.

Based on our assumptions about the market price of electricity and the curtailment cost of the customer, our model makes several practical recommendations. EnerNOC would be better off introducing a participation fee to extract extra surplus from Four Seasons instead of a no-fixed payment contract. For low levels of cost uncertainty, an automated contract with a two-part tariff is EnerNOC’s most profitable contract choice. For high cost uncertainty, EnerNOC earns a higher profit by offering a voluntary contract. Note that a simple linear voluntary contract is relatively close to optimal in this region.

Furthermore, our experiment suggests that firm and social planner incentives are aligned when the variance of cost is either small or large. When the standard deviation of cost is roughly between $[4600, 8200] \$/MW^2$, the voluntary contract creates more welfare for the system, but EnerNOC prefers the automated one. In this region, the automated contract allows the firm to extract more surplus, therefore the social planner and the firm incentives are not aligned.

1.8. Conclusion

In this paper we introduce the voluntary contract to the domain of electricity curtailment contracts and compare its performance with the conventional automated contract. We

assume that a curtailing consumer's load reduction is limited by a stochastic curtailment cost, which is unobserved until the curtailment event. The consumer derives value from the voluntary contract which allows the consumer to observe his curtailment cost before committing to a particular load reduction level. In contrast, the automated contract requires the consumer to commit to a load reduction level without knowledge of his cost. Critically, the firm is punished for suboptimal load delivery through changes in the market price. When deciding between these contracts the firm faces a trade off between the cost of procuring load and the price impact of the load reduction.

Comparing the performance of these two contracts we find that a firm's and a social planner's preferences are driven both by market conditions and consumer characteristics. In markets with prices that are highly sensitive to supply, the certainty of the level of curtailed load delivered under the automated contract is more valuable to the firm than the reduction in the payments made to consumers under the voluntary contract. This difference in firm profits is also displayed in social welfare. It follows that firms serving as emergency reserves, when the market price curve is at its steepest, should offer automated contracts exclusively. Consumers also determine a firm's contract choice through their cost profiles. If market price sensitivity is not too high and the curtailment cost distribution has a large Jensen's gaps, the voluntary contract yields higher firm profit and social welfare. This result distills the contract choice intuition into a single attribute of the consumer's cost distribution. Furthermore, we show that, for symmetric cost distributions, a firm can use the variance of cost to distinguish between customers that should be offered voluntary or automated contracts: higher variance favors voluntary contracts. As curtailment contracts become more broadly used we hope that these findings will serve as a guide for practitioners to tailor their programs to their market and their customers.

These results build on the literature studying firms that allow flexibility to autonomous suppliers (in this paper, consumers supply curtailment to the firm). In our model, the price-taking firm faces consumers with perfectly correlated cost of curtailment. This correlation

causes the firm to bear price risk, or face uncertain market price caused by the variable curtailment allowed by the voluntary contract. When the market price is sensitive and when consumers derive little value from the flexibility of the voluntary contract, the cost of this price risk outweighs the benefits of flexibility, and the firm prefers the automated contract.

This work has practical implications for regulators in the field of electricity demand response. We show that, from an environmental perspective, a deregulated market for demand response leads to the ideal contract choice (maximum electricity consumption is avoided). On the other hand, considering the economic surplus of the customer and the curtailment service provider, the firm does not always choose the welfare maximizing contract. A social planner could improve the welfare of the system by encouraging voluntary contracts in some cases in which low cost volatility or high market price sensitivity cause the firm to prefer the automated contract.

The authors thank Eduardo Azevedo, Volodymyr Babich, Gerard Cachon, Morris Cohen, Noah Gans, Daniel Gottlieb, Ilan Lobel, Kenneth Ogawa, Sergei Savin, and Xuanming Su for their thoughtful comments. We also thank the Wharton Risk Management Centers Russell Ackoff Doctoral Student Fellowship for the generous financial support.

CHAPTER 2 : The Role of Surge Pricing on a Service Platform with Self-Scheduling Capacity

2.1. Introduction

¹ The rise of the “sharing economy” has transformed the way firms can deliver service to consumers. The firm no longer must centrally schedule its capacity by assigning workers to shifts. Instead, workers may act as independent service providers who determine their own work schedules, and the firm’s role becomes that of a platform which connects providers to consumers. (See Katz and Krueger (2016) for data on the growth of alternative work arrangements in the United States.) Although the platform has far less control over how many providers work at any one time, providers gain the freedom to “self-schedule” the hours they work, presumably allowing them to better integrate their work with the other activities in their lives (Hall and Krueger (2015b)). To make these new relationships viable, customers must be charged a reasonable fee and be adequately served.

Examples of relatively new platforms that feature self-scheduling capacity include Uber and Lyft for local transportation, and Postmates and Instacart for local delivery. A potential provider for one of these platforms must first make the long-term decision of whether to join the platform or not. This decision has implications for several months or years, and providers join only if they expect to earn more with the platform than with their next best alternative. If a person joins a platform as a provider, then they must make short-term decisions about when and how often to work. These decisions are made on a daily or hourly basis, so the participation decision is relevant over a much shorter time interval than the joining decision. The participation decision is based in part on the wage providers receive per service. It is also based on providers’ expectations of how likely they are to get work, which is a function of the overall level of demand and the number of providers working at that time on the platform. For example, an Uber driver may know that demand is higher on

¹Reproduced with permission. Copyright, INFORMS, <http://www.informs.org>

rainy days but may also know that other drivers are more likely to drive as a consequence. What matters to the provider is the amount of demand relative to the amount of offered capacity at a particular time.

In this paper we focus on the contractual forms a monopoly platform could select to make a viable market with self-scheduling capacity. We study a stylized model with the following features: (i) there exists a large pool of potential providers, (ii) providers join the platform only if their rational expectation of their earnings from participation on the platform exceeds the stochastic opportunity cost of their next best activity, (iii) the platform sets a price for consumers, a wage paid to providers for work completed and regulates the maximum number of providers who join the platform, (iv) the platform cannot directly determine when providers work and, instead, the providers who joined the platform self-schedule their offered capacity, (v) demand is stationary but varies in predictable ways (e.g., more consumers seek transportation on a rainy evening), (vi) if the offered capacity exceeds demand, providers share the available demand equally, but if the offered capacity is less than demand, then demand is randomly rationed (i.e. all consumers are equally likely to receive the scarce service), (vii) the platform's price and wage can depend on the current level of demand and (viii) provider's opportunity costs are independent and identically distributed across providers and time.

There are three key features of the model that make this environment distinctive and capture some of the interesting dynamics of these service platforms in practice. First, providers self-schedule their offered capacity. Consequently, even if the number of providers who have joined the platform is sufficient to satisfy demand, it is possible that either *demand rationing* (too few providers choose to work) or *capacity rationing* (too many providers choose to work) can occur. Both forms of rationing represent costly inefficiencies for the platform. Second, the platform can offer demand-contingent prices and wages. Demand-contingent prices are often called *dynamic prices*. Uber and Lyft employ versions of dynamic prices and wages called *surge pricing* and *prime time* respectively. There is a large literature on dynamic

prices, while the literature on dynamic wages is far less extensive, and there is no work on the interaction between dynamic prices and dynamic wages. Third, capacity decisions are made at two different time scales: providers make a “long run” decision to join the platform or not and then in the “short term” decide whether to participate or not. At the time the participation decision is made, the joining decision (and cost) is sunk.

The platform’s primary goal with the design of its contract is to maximize its profit. Doing so requires a contract that assures providers that join sufficient expected profit. However, the contract must not give providers too much of an incentive to participate, which could lead to an excess of providers, nor too little incentive, which could entice too little participation from providers to satisfy demand.

Although maximizing profit is a clear objective for the platform, it is not the platform’s only concern. A number of controversies have emerged with this new business model. Some people believe providers are not adequately compensated because they are not given benefits and rights associated with being employees (Isaac and Singer (2015), Scheiber (2015)). Others worry that customers are unfairly discriminated against as a result of dynamic pricing (Kosoff (2015), Stoller (2014)). Consequently, with a view towards potential litigation and regulation, a platform should be concerned with both provider and consumer welfare. In particular, it is important to understand the degree to which there is a tension between maximizing the platform’s profit and the surplus earned by the other relevant stakeholders, the providers and consumers.

We focus on five possible operating models, or contracts, for the platform. With the simplest possible contract, called the *fixed contract*, the platform offers providers a fixed wage and charges consumers a fixed price. Next, we consider contracts in which the the platform either chooses dynamic prices (with a fixed wage), or dynamic wages (with a fixed price). We refer to the former as the *dynamic price contract* and the latter as the *dynamic wage contract*. A *commission contract*, which resembles surge pricing used in practice, allows the platform to dynamically adjust both prices and wages in response to demand, but imposes

the constraint of a fixed commission, i.e., a fixed ratio between the two. The commission contract is used in practice; for example Uber offers its drivers a fixed 80% commission in most markets (Huet (2015)). It has been argued that this constraint may substantially lower the platform's profit (Economist (2014)). Finally, the platform's *optimal contract* dynamically adjusts both prices and wages without the constraint of a fixed commission. A closed form solution for the best version of each of these contracts is unavailable, but we analytically determine how to determine the best form of each contract with a single dimensional search over a bounded space. In addition, we are able to analytically determine conditions under which a commission contract is optimal for the platform. Via numerical analysis over the set of feasible and plausible parameters, we compare profits, consumer surplus and provider surplus across all five contracts. Those results are consistent with the analytical results derived from a special case of the model.

To preview our main results, we find that the optimal contract provides the platform substantially higher profit relative to the fixed contract and self-scheduling is a profitable arrangement for the platform relative to central-scheduling. Although not optimal, the commission contract is nearly optimal, and given its simplicity, this may explain its use in practice. We find that consumers indeed have a reason to be skeptical about dynamic pricing: relative to the fixed contract, adding dynamic pricing (with a fixed wage) reduces consumer surplus - the platform uses dynamic pricing to extract consumer surplus for its own profit. However, again relative to the fixed contract, adding dynamic pricing and dynamic wages together can increase consumer surplus even though that combination also maximizes the platform's profit - the added value created by reducing capacity and demand rationing allows all parties to be better off. It does so when the fixed contract rations demand when demand is high, which is when demand rationing due to limited capacity is particularly costly. Thus, if the lack of dynamic prices and wages leads to poor service for customers in high demand periods, then consumers actually benefit from the introduction of dynamic pricing, like Uber's surge pricing.

2.2. Literature Review

Our work is primarily connected to three domains in the existing literature: research on capacity and pricing, revenue management models, and recent papers on peer-to-peer platforms and self-scheduling capacity. For simplicity and consistency, we refer to the various components in other papers using the terms relevant for our model. For example, the “platform” is the organization responsible for designing the market, “providers” generate capacity, “dynamic prices” are demand-contingent payments from consumers to the platform in exchange for service, and “dynamic wages” are demand-contingent payments from the platform to providers.

Several papers study competition among multiple providers and establish that competition can lead to excessive entry (e.g. Mankiw and Whinston (1986)) and a platform should discourage competition to mitigate the losses in system value due to this issue (e.g. Bernstein and Federgruen (2005), Cachon and Lariviere (2005)), but those papers do not consider dynamic wages or prices.

A set of papers considers peak-load pricing, the practice of charging higher prices during peak periods of demand (e.g. Gale and Holmes (1993)). The primary motivation of peak-load pricing is to increase revenue by shifting demand from the peak period to the off-peak period. We do not incorporate this capability into our model. For example, consumers in need of transportation during a rainy evening are unable to postpone their need to a time with better weather.

There is work on the value of dynamic prices in systems that experience congestion, but with fixed capacity: e.g., Celik and Maglaras (2008), Ata and Olsen (2009), and Kim and Randhawa (2015). Banerjee, Johari, and Riquelme (2015) considers the value of dynamic pricing in a model with random arrivals of consumers and providers. Unlike us, they find that dynamic pricing provides no benefit in terms of maximizing the platform’s expected profit or system welfare, but they have a single demand regime whereas in our model some

periods (importantly) have predictably higher demand than others for a given price.

There is a considerable literature on “two-sided markets” in which platforms earn rents by creating a market for buyers and sellers to transact (e.g., a game console maker as the platform between game developers and consumers). These papers tend to focus on which side of the market the platform charges based on the various externalities within the system but they do not consider dynamic demand (e.g., Rochet and Tirole (2006)).

Peer-to-peer service platforms have attracted significant academic interest; e.g. Kabra, Belavina, and Girotra (2015), Hong and Pavlou (2014), Snir and Hitt (2003), Moreno and Terwiesch (2014), and Yoganarasimhan (2013). Those papers investigate how to subsidize different market players to accelerate the growth of a peer-to-peer platform, whether consumers have geographic preferences over providers, the influence of platform design on provider quality, and how provider reputation impacts the market. We do not explore those issues: our providers are ex-ante homogeneous and do not build reputations. Fraiberger and Sundararajan (2015) investigate the interaction between ownership and sharing on a peer-to-peer marketplace, a dynamic that is not addressed in our model. Cohen et al. (2016) use Uber transaction data to measure the amount of consumer surplus generated given the implementation of surge pricing, but they do not estimate a counterfactual consumer surplus level with other contractual forms.

There is modeling and empirical work on the competition between peer-to-peer service marketplaces and existing markets: Einav, Farronato, and Levin (2015), Zervas, Proserpio, and Byers (2016), Seamans and Zhu (2013), Cramer and Krueger (2016), and Kroft and Pope (2014). We do not directly consider the competition between the platform and incumbents.

Several papers (e.g. Hu and Zhou (2015) and Allon, Bassamboo, and Çil (2012)) explore the process for matching providers to consumers when capacities are exogenous and all participants have preferences for the match they receive (e.g. a courier prefers to be matched to a nearby consumer). We do not consider matching because our consumers and providers

are homogeneous, so careful matching does not provide a benefit.

Closest to our work are papers on self-scheduling capacity. Ibrahim and Arifoglu (2015) considers a model in which the platform chooses the number of providers and providers are either assigned by the platform to work in one of two different periods or they self select which of the two periods they work in. Unlike in our model, the platform can directly control the number of providers in the system. Taylor (2016) and Bai et al. (2016) study queuing systems in which a platform creates a market for service where arrivals of consumers and servers are endogenously determined based on decisions to seek and provide service respectively. Their models do not consider dynamic prices or wages, and the number of potential providers is exogenous (i.e., capacity decisions are made on a single, short-term, time scale). Gurvich, Lariviere, and Moreno (2015) studies a model in which a platform directly chooses the number of available providers, the wage for each provider who chooses to work, and a cap on the number of providers who are allowed to work: given the platform’s prevailing wage, more providers may want to work than the platform wants. They do not include dynamic pricing - in all versions of their model the platform selects a single price. They also do not impose an earnings constraint for providers. Instead, they impose an exogenous minimum wage. In our model providers decide whether to join the platform based on rational expectations of future earnings.

2.3. Model




Figure 5: Timeline of events

As shown in Figure 5, the interaction between the platform, providers, and consumers occurs over two stages, or periods. At the start of the first period the platform announces the terms of trade, consisting of prices charged to consumers, wages paid to providers, and the maximum number of providers allowed to join the platform. A large pool of potential

providers then decides whether to join the platform or not. We refer to this as the “joining” decision. This period represents the providers’ long-term decision. With a ride-sharing platform such as Uber, period 1 would represent a provider’s decision to sign up for Uber instead of Postmates, for example. The second period represents the short-term decisions to work on the platform or not. We refer to this as the “participation” decision. For example, once on the platform, providers for Uber must decide whether to offer their service during a particular day or even a particular hour. Consequently, the participation decision is relevant over a much shorter time interval than the joining decision. Hence, the provider expects to make many of these short-term decisions. For simplicity, we collapse these decisions into a single period.

In period 1 a provider incurs an opportunity cost, c_1 , for joining the platform and in period 2 the provider can earn a profit from participation on the platform. Hence a provider joins in period 1 only if the provider expects to earn in period 2 at least c_1 . All providers share the same opportunity cost, so either all are willing to join or none are. Our model approximates a market with a deep pool of potential providers and a highly elastic supply curve: if expected earnings are less than c_1 , then the number of interested providers drops substantially, but if greater than c_1 , then there is an ample number of interested providers.

There are two types of uncertainty. The first is each provider’s cost to participate on the platform in period 2. For example, on some days participation might be costly (e.g. a child needs to visit a doctor) while on other days participation isn’t costly (the provider has nothing else to do that day). Each provider can anticipate in period 1 that they will incur a participation cost in period 2, but they do not know what that cost will be. They learn their participation cost at the start of period 2 before their participation decision. In particular, let c_2 be a provider’s realized participation cost in period 2. The stochastic participation cost is independently and identically distributed across providers with distribution $G(c)$ and density $g(c)$, which are known at the start of period 1 at the time of the provider’s joining decision. We assume $G(c)$ is strictly increasing and differentiable, $G(0) = 0$, and there does

not exist a finite c such that $G(c) = 1$. In section 2.5 we consider a simplified version of the model in which the participation is fixed across all providers.

The demand level is the second type of uncertainty. Demand occurs only in period 2 and it can be either “high” or “low”. For example, for a ride-sharing platform, “high” demand could be a rainy evening on a holiday weekend, whereas “low” demand could be a warm Wednesday evening. The platform and the providers can anticipate in period 1 that demand can be either high or low, but they only learn the actual state of demand at the start of period 2, after their joining decision but before their participation decision. Thus, providers make their joining decision before either uncertainty is resolved, but they make their participation decision after observing both demand and participation cost. Note, while each provider observes their own c_2 , the platform does not observe each provider’s participation cost, so only demand uncertainty is resolved for the platform.

The platform faces a linear demand curve with an uncertain intercept. To be specific, demand for the platform’s service is $D_j = (a_j - bp_j)^+$, where p_j is the price charged to consumers, b is a constant, and the demand state can either be low or high, $a_j \in \{a_l, a_h\}$, where $a_l < a_h$. Let f_j , $j \in \{l, h\}$ be the probability of state j demand, where $f_l + f_h = 1$. Each participating provider can serve up to a single unit of demand in period 2. The parameter b has no impact on the qualitative results, so $b = 1$ is assumed throughout.

At the start of period 1 the platform announces the terms of trade for providers joining the platform. The terms consist of (i) an upper bound, N , on the number of providers who can join (e.g. Uber imposes a cap on the total number of drivers that can operate in a city), (ii) a price charged to consumers in each demand state, p_j , and (iii) a wage paid in each demand state to each provider for service, w_j . We say that the platform uses demand-contingent, or dynamic, prices if $p_l \neq p_h$. The platform can also choose a single price no matter the demand state, i.e. $p_l = p_h$. The same applies for wages.

For a particular demand realization, price, and wage, it is possible that demand exceeds

the capacity of participating providers. In that case demand is randomly rationed: some demand is not served while all participating providers serve one unit of demand. Alternatively, it is possible that capacity exceeds demand. In that case capacity is rationed: participating providers utilizes only a portion of their capacity. To be specific, let ϕ_j be a provider's utilization in demand state a_j , where ϕ_j is the fraction of capacity offered by the participating providers used to serve demand. When demand is rationed, $\phi_j = 1$, whereas when capacity is rationed, $\phi_j < 1$.

A participating provider earns revenue $\phi_j w_j$ in period 2. All providers (who joined in period 1) with participation cost $\phi_j w_j$ or lower choose to participate, while providers unfortunate to have high participation costs choose not to participate. We require that providers make maximizing decisions based on rational expectation regarding their earnings. (See Farber (2015) and Chen and Sheldon (2015) for evidence that taxi drivers and Uber providers respectively make decisions based on rational expectations to maximize their return.) Thus, assuming N providers join the platform in period 1, in equilibrium

$$\phi_j = \begin{cases} 1 & NG(w_j) \leq a_j - p_j \\ \frac{a_j - p_j}{NG(\phi_j w_j)} & a_j - p_j \leq NG(w_j) \end{cases}$$

Note that in the second case with capacity rationing, i.e. $a_j - p_j \leq NG(w_j)$, a recursive relationship determines the equilibrium utilization. This equilibrium utilization exists and is unique.

Let π_j be a provider's expected profit conditional on joining for a given demand state a_j , wage w_j , and price p_j :

$$\pi_j = (w_j \phi_j - E_{c_2}[c_2 | c_2 \leq w_j \phi_j])G(w_j \phi_j) = \int_0^{w_j \phi_j} G(c)dc$$

Let Π be a provider's expected profit from joining the platform:

$$\Pi(p, w, N) = \sum_{j \in \{l, h\}} \left(\int_0^{w_j \phi_j} G(c) dc \right) f_j$$

If $c_1 \leq \Pi(p, w, N)$, then all potential providers attempt to join the platform, but the platform's imposed cap of N limits the number that actually join to the N . However, if $\Pi(p, w, N) \leq c_1$, then no providers join. Hence, for the platform to function, it must offer terms such that $c_1 \leq \Pi(p, w, N)$. Throughout we assume that such terms are offered and hence N providers join the platform.

The platform's objective is to choose price, wage, and recruitment to maximize its expected profit subject to the (already mentioned) constraint that providers are willing to join the platform:

$$\begin{aligned} \max_{w, p, N} U(p, w, N) &= \sum_{j \in \{l, h\}} (p_j - w_j) \phi_j N G(\phi_j w_j) f_j \\ \text{s.t. } c_1 &\leq \Pi(p, w, N) \end{aligned}$$

It is helpful for our analysis to implicitly define four parameters, w' , w'' , $\bar{\phi}_l$, and \bar{c}_1 :

$$\int_0^{w'} G(c) dc = c_1; \quad \int_0^{w''} G(c) dc f_h = c_1; \quad \int_0^{\bar{\phi}_l w} G(c) dc f_l + \int_0^w G(c) dc f_h = c_1; \quad \bar{c}_1 = \sum_{j \in \{l, h\}} \int_0^{a_j} G(c) dc f_j$$

The first, w' , is the smallest wage that induces providers to join when they can assume that they are assured to be paid w' in either demand state in equilibrium. The second, w'' , is similar to w' , except this is the lowest wage that induces providers to join when they are assured to receive w'' payment in the high demand state and no payment in the low demand state. (If $a_l \leq p$, then there are no customers to serve in the low demand state.) The third, $\bar{\phi}_l$, which applies when $w' < w < w''$, is the rational expectations equilibrium utilization when providers expect to be rationed in the low demand state but not in the high demand state. The fourth, \bar{c}_1 , is the maximum joining cost that allows for a positive surplus in the

system (i.e., if $\bar{c}_1 < c_1$ then a provider wouldn't join the platform even if she were the only provider on the platform and the platform allowed her to keep all of the possible profit). As $\bar{c}_1 < c_1$ means this market cannot function, we assume $c_1 < \bar{c}_1$ throughout.

Beside its own profit, the platform may have an interest in consumer and provider surplus, especially if the platform's practices are potentially controversial, thereby motivating negative publicity, lawsuits, or government regulation. We measure consumer surplus under a linear stochastic demand in a similar fashion to Cohen, Lobel, and Perakis (2015): $S = \sum_{j \in \{l, h\}} 0.5 \min((a_j - p_j)^2, (a_j - p_j)NG(\phi_j w_j))f(a_j)$. Consumer surplus decreases in the prices charged and increases in the number of consumers served. The latter depends on the number of providers that join the platform, N , and the fraction of those recruited providers that decide to participate. Provider surplus increases in the number of recruited providers and in those providers' expected earnings. If each provider earns exactly c_1 conditional on joining (as is shown in each of the contracts we consider), then total provider surplus is $c_1 N$.

2.4. Contract Design

We focus on five contract designs that vary by the amount of flexibility given to the platform to adjust its prices and wages in response to observed demand in period 2. A closed form solution for the platform's best version of each contract is unavailable, but the following five theorems indicate that the platform's best contract within each design can be found via a single dimensional search over a bounded interval (even though each contract involves up to five decisions: a price and wage for each demand state and the number of providers to allow on the platform). Proofs are available in the appendix.

2.4.1. Fixed Contract

With the fixed contract the platform chooses a single per-service wage, w , to pay providers and a single per-service price, p , to charge consumers. These quantities are independent of the realized demand state. As a result, the platform is subject potentially to two inefficien-

cies: demand rationing and capacity rationing. With demand rationing, the offered wage is too low to induce enough providers to participate relative to realized demand, leaving some customers without service. With capacity rationing, the offered wage is too high because too many providers participate relative to realized demand. For a given contract, it is possible that demand is rationed in the high demand state and capacity is rationed in the low demand state, as is illustrated in Figure 6. In the low demand state, $NG(\phi_l w)$ providers participate, which exceeds demand, $D_l = a_l - p$. In the high demand state, $NG(w)$ providers participate, all are allocated a customer, but $D_h - NG(w)$ customers do not receive service, even though the number of providers on the platform, N , may be adequate to serve all demand.




Figure 6: An example of demand and capacity rationing with a fixed contract.

The fixed contract may not be able to earn a positive profit (given $c_1 < \bar{c}_1$), but if it does so, then Theorem 6 describes the best fixed contract for the platform, which can be divided into two types: (i) the platform serves both demand states, or (ii) the platform only serves high demand. There are two extreme versions of serving demand in both states. In the first, which we refer to as the *poor service* version, capacity matches low demand, meaning that there is no capacity rationing and providers are fully utilized in all states. However, while all customers are served in the low demand state, in the high demand state $a_h - a_l$ of demand is lost. In the second version, which we refer to as the *poor utilization* version, capacity matches high demand. Customers are fully served in either state, but in the low demand state too many providers participate, chasing too little demand, leading to capacity rationing.

Theorem 6. *Conditional on earning a positive profit, the best fixed contract has one of the following two characteristics:*

1. The platform serves both demand states. In particular, $w \in [w', \min(a_l, w'')]$,

$$p = \max\left(\frac{a_l + w}{2}, \frac{G(w)a_l - \bar{\phi}_l G(\bar{\phi}_l w)a_h}{G(w) - \bar{\phi}_l G(\bar{\phi}_l w)}\right)$$

there is demand rationing only in the high state (i.e. $N = (a_l - p)/(\phi_l G(\phi_l w))$), there is capacity rationing only in the low state (i.e. $\phi_l = \bar{\phi}_l \leq 1$ and $\phi_h = 1$), and each provider's joining constraint is binding, i.e. $c_1 = \Pi(p, w, N)$.

2. The platform serves only high demand. In particular, $w = \min\{w'', a_h\}$, $a_l < p = (a_h + w)/2$, $N = (a_h - p)/G(w)$, and participating providers are fully utilized, i.e. $\phi_h = 1$.

2.4.2. Dynamic Wage Contract

With the dynamic wage contract the platform charges consumers a fixed price, p , but pays providers a wage, w_j , that depends on the demand state a_j . Relative to the fixed contract, the dynamic wage contract allows the platform to address the issue of capacity rationing due to excessive provider participation. For example, suppose the platform's fixed contract rations capacity in the low demand state. The platform could lower its wage in the low demand state while leaving providers no worse off; providers would be paid less but, because fewer providers participate, their utilization would increase. Consequently, the platform's profit would strictly increase. Alternatively, suppose the platform's fixed contract rations demand in the high demand state. This is the best fixed contract when it is too costly to regulate provider participation with a single wage, so it is regulated by restricting recruitment in the first stage, N . However, because a demand-contingent wage gives the platform a greater ability to regulate provider participation, the platform may no longer need to rely exclusively on restricting recruitment, allowing higher N , thereby mitigating some demand rationing. In fact, according to Theorem 7, the dynamic wage contract is capable of eliminating capacity rationing in all demand states. However, the best dynamic wage contract may still ration demand, which is why it may not be able to earn a positive profit.

Theorem 7. *Conditional on earning a positive profit, the best dynamic wage contract has one of the following two characteristics:*

1. *The platform serves both demand states. In particular,*

$$c_1 = \int_0^{w_l} G(c)dcf_l + \int_0^{w_h} G(c)dcf_h$$

$$p = \max \left(\frac{a_h G(w_l) - a_l G(w_h)}{G(w_l) - G(w_h)}, \min \left(a_l, \frac{a_l}{2} + \frac{f_h G(w_h)w_h + f_l G(w_l)w_l}{2(G(w_h)f_h + G(w_l)f_l)} \right) \right)$$

there is demand rationing only in the high state (i.e. $N = (a_l - p)/G(w_l)$), there is no capacity rationing, i.e. $\phi_l = \phi_h = 1$, and each provider's joining constraint is binding, i.e. $c_1 = \Pi(p, w, N)$.

2. *The platform serves only high demand. In particular, $w_h = \min \{w'', a_h\}$, $p = (a_h + w_h)/w_h$, $N = (a_h - p)/G(w_h)$, and participating providers are fully utilized, i.e. $\phi_h = 1$.*

2.4.3. Dynamic Price Contract

With the dynamic price contract, the platform selects a price for each demand state, p_j , but pays providers a fixed wage. The dynamic price contract enables the platform to manage demand rationing. For example, suppose the best fixed contract has poor service. Capacity is restrictive because higher capacity would lead to costly capacity rationing in the low demand state. However, with dynamic prices the platform can increase its price in the high demand state without affecting providers, thereby reducing demand rationing while increasing its revenue and profit. With the other extreme, suppose the best fixed contract has poor utilization. In the high demand state, the platform would prefer to raise the price further. But doing so would exacerbate the problem of capacity rationing in the low demand state. Once the platform has the ability to charge dynamic prices, it can indeed raise its price in the high demand state while also lowering its price in the low demand state, both of which help to mitigate capacity rationing while still avoiding demand

rationing. Nevertheless, a positive profit is not always feasible.

Theorem 8. *Conditional on earning a positive profit, the best dynamic price contract has one of the following two characteristics:*

1. *The platform serves both demand states. In particular, $w \in [w', \min(a_l, w'')]$, $p_l = (a_l + w)/2$, $p_h = a_h - G(w)N$, $N = (a_l - w)/(2\bar{\phi}_l G(\bar{\phi}_l w))$, there is no demand rationing, there is capacity rationing only in the low state, i.e. $\phi_l = \bar{\phi}_l \leq 1$, and $\phi_h = 1$, and each provider's joining constraint is binding, i.e. $c_1 = \Pi(p, w, N)$.*

2. *The platform serves only high demand. In particular, $w = \min\{w'', a_h\}$, $p = (a_h + w)/w$, $N = (a_h - p)/G(w)$, and participating providers are fully utilized, i.e. $\phi_h = 1$.*

2.4.4. Commission Contract

The commission contract, which resembles Uber's surge pricing policy, adjusts both price and wage in response to demand, but also imposes the constraint that the two have a constant ratio. In particular, the platform charges a demand-contingent price, p_j , and pays providers $w_j = \beta p_j$, where β is the (fixed) commission rate. Given the market is viable ($c_1 < \bar{c}_1$), there exists a sufficiently high commission rate that enables the market to function and the platform to earn some profit.

For a given commission, there is a unique best wage schedule and recruitment level satisfying the optimality conditions in the following theorem, but a line search is required to find the best commission.

Theorem 9. *For a given $\beta \in \left[\frac{w'}{a_h}, 1\right]$, the best fixed commission contract is uniquely defined, earns a positive profit for the platform and satisfies:*

$$p_j = \max\left\{a_j - NG(\hat{w}_j), \frac{1}{2}a_j\right\}; \quad \phi_j = \min\left(1, \frac{a_j}{2NG(\frac{1}{2}\beta\phi_j a_j)}\right); \quad c_1 = \sum_{j \in \{l, h\}} \int_0^{w_j \phi_j} G(c) dc f_j$$

where \hat{w}_j is uniquely defined by $\hat{w}_j = \beta(a_j - NG(\hat{w}_j))$. The providers' joining constraint is binding. Capacity rationing is possible, but demand rationing does not occur.

2.4.5. Optimal Contract

The optimal contract allows the platform complete flexibility: both wages and prices may vary according to the demand state without the constraint of a fixed ratio between the two. With these two levers, the platform maximizes its profit, it eliminates both demand and capacity rationing, it always serves demand in all demand states, and it maximizes system surplus (the sum of platform and provider expected profits).

Theorem 10. (i) *The platform earns a positive profit with the optimal contract (for all $c_1 < \bar{c}_1$), (ii) the optimal contract is uniquely defined by w , p , and N satisfying,*

$$w_j = a_j - 2NG(w_j); \quad p_j = a_j - NG(w_j); \quad c_1 = \sum_{j \in \{l, h\}} \int_0^{w_j} G(c) dc f_j$$

iii) *there is no capacity rationing, i.e. $\phi_l = \phi_h = 1$, nor demand rationing, (iv) each provider's joining constraint is binding, i.e. $c_1 = \Pi(p, w, N)$, and (v) system surplus (the sum of platform and provider profits) is maximized.*

For a given N , the system of the first two equations uniquely identifies prices and wages. A search over N finds the contract that satisfies all three equations.

Unlike the commission contract, the optimal contract is not burdened with the constraint of a fixed ratio between wage and price. Nevertheless, there are cases in which the optimal contract is a commission contract (i.e., the commission contract is optimal for the platform). For example, the optimal wage to price ratio, $(w_j/p_j) = w_j/(w_j + NG(w_j))$, is independent of the demand state (i.e., constant across states) if participation costs are uniformly distributed (i.e., $G(c)$ is linear in c). Alternatively according to Theorem 11, the commission contract is optimal if joining costs are either very low or very high. To explain, when the joining cost, c_1 , approaches its upper bound \bar{c}_1 , the optimal contract gives nearly all revenue to providers to recruit them. This is equivalent to a commission contract with $\beta \rightarrow 1$. When c_1 instead approaches zero, the platform can recruit many providers and encourage enough participation with a very small wage. In the limit, the optimal contract

offers almost no wages, which is equivalent to a commission contract with $\beta \rightarrow 0$.

Theorem 11. *The commission contract is optimal (i.e., yields the same profit for the platform as the optimal contract) if (i) $c_1 \rightarrow \bar{c}_1$ or (ii) $c_1 \rightarrow 0$.*

2.5. Fixed Participation Cost

In this section we consider a specialized version of the main model in which, instead of heterogeneous and stochastic participation costs with infinite support described by the distribution function $G(\cdot)$, all providers have a fixed participation cost, c_2 , in period 2. (i.e., $G(c|c < c_2) = 0$ and $G(c|c_2 \leq c) = 1$.) All other aspects of the main model remain. Hence, this *fixed* c_2 model, retains most of the critical features of the main model: e.g., providers act on rational expectations, capacity and demand rationing are possible, and supply decisions are made over two time scales.

To conserve space, we focus on three contract types with the fixed c_2 model: (1) a fixed contract, (2) the optimal contract (i.e., dynamic prices and wages), and (3) the commission contract (i.e., dynamic prices and a fixed ratio between wage and price). With the fixed contract the platform selects a fixed price and compensates the providers so that their joining constraint binds, i.e., they each earn c_1 . Hence, the fixed contract in this model is comparable to the fixed contract in the main model.² For notational convenience, let $\bar{a} = f_l a_l + f_h a_h$ and $\hat{c} = c_2 + c_1/f_h$. See the online appendix for proofs and derivations of results.

The primary objective of the fixed c_2 model is to use its additional tractability to derive analytically (i) the conditions under which the optimal contract increases consumer surplus relative to the fixed contract, and (ii) a lower bound for the platform's profit with the commission contract relative to the optimal contract. The numerical calculations in the subsequent section demonstrate that these results carry over to the (more general) main

²This compensation can be achieved with a fixed wage for service (equal to c_2 , so that all providers who participate receive demand) and a fixed salary for joining the platform (equal to c_1 , to ensure the joining constraint is satisfied).

model.

In the fixed c_2 model, the best fixed contract adopts one of three possible versions: (i) a “poor service” version with demand rationing; (ii) a “poor utilization” version with capacity rationing; (iii) a “only high demand” version in which no demand is served in the low demand state. The optimal contract, serves both demand states and sets recruitment, N , equal to high demand.

Proposition 1 identifies the situations in which the optimal contract increases consumer surplus relative to the fixed contract. If providers are relatively expensive (high c_1) then the fixed contract involves demand rationing (poor service) and consumers benefit from switching from the fixed contract to the optimal contract. In these cases the fixed contract is unable to provide adequate supply and even though consumers pay more in the high demand state with the optimal contract, the additional supply available with the optimal contract leads to higher consumer surplus. However, if providers are relatively cheap (low c_1) then the fixed contract rations capacity (e.g., the poor utilization version), and consumers are worse off with a switch to the optimal contract.

Proposition 1. *In the fixed c_2 model, the optimal contract has higher consumer surplus than the fixed contract if and only if “poor service” or “only high demand” is the best version of the fixed contract.*

The commission contract is the third contract of interest. There are three versions of the commission contract - three of them yield closed form solutions whereas the fourth does not. The fourth version is not problematic for two reasons - it is the least likely of the versions to be the best commission contract, and it is not necessary to include in the derivation of the lower bound profit ratio in Proposition 2.)

The optimal contract is a commission contract when the joining cost is sufficiently high: if $f_h(a_h - a_l) < c_1$ then the optimal contract chooses the same commission in either demand state, so a commission contract with a single commission can replicate the optimal contract.

In contrast, if the joining cost is “low” (i.e., $c_1 \leq f_h(a_h - a_l)$), then the optimal contract chooses commission rates that differ across the demand states, i.e., $\beta_l = w_l/p_l \neq w_h/p_h = \beta_h$. In these cases the commission contract must select a commission rate that is sub-optimal in one or both states, reducing the platform’s profit with the commission contract relative to the optimal contract.

Proposition 2. *The following is a lower bound for the ratio of the platform’s profit with the commission contract, U_β , and the platform’s profit with the optimal contract, U_o : $\min\{U_\beta/U_o\} = (1 + \sqrt{f_h})/2$. This bound is achieved either when $c_1 = 0$ or $c_2 = 0$.*

Proposition 2 reports on a lower bound for the platform’s profit with the commission contract. The commission contract performs poorly when one of the two costs is very low (either c_1 or c_2) and the probability of high demand is small. In the extreme, as $f_h \rightarrow 0$, the fixed commission contract earns only 1/2 of the profit of the optimal contract. However, when the two demand states are equally likely, the commission contract earns at least 85% of the optimal profit $\left((1/2) \left(1 + \sqrt{1/2}\right)\right)$. As $c_2 \rightarrow 0$, the optimal contract chooses a low commission when demand is low (to prevent too much participation) and, when demand is high, chooses a sufficiently high commission to give providers enough profit (c_1/f_h) to justify joining the platform. This disparity in the two commissions creates a challenge for the commission contract, which is required to choose a single commission. With the other extreme, $c_1 \rightarrow 0$, the joining constraint is not important. Instead, the focus is on the incentive for providers to participate. Because $p_l < p_h$, which implies $c_2/p_h < c_2/p_l$, the best commission with low demand is higher than with high demand (because both states must yield at least c_2 for the providers to participate). Again, the commission contract does not do well with this disparity in commissions. Note, according to Theorem 11, the commission contract yields the optimal profit as $c_1 \rightarrow 0$ in the main model, which contrasts sharply with its performance in the fixed c_2 model. The difference occurs because in the fixed c_2 model $G(\cdot)$ has finite support whereas in the main model it has infinite support. Consequently, in the fixed c_2 model the average participation cost conditional on participation is independent of the number of joining providers, N , (i.e., it is always c_2), whereas in the main model it

decreases in N (i.e., for the same desired number of participating providers, increasing N lowers the average participation cost).

Although there are cases in which the commission contract performs poorly relative to the optimal contract, this does require special parameters. For example, consider only the extreme cases in which $f_h = 0.05$, which yields a lower bound of $U_\beta/U_o = 0.612$. Evaluation of 3,600 evenly spaced observations throughout the feasible parameter space yields a minimum profit ratio close to the lower bound, $U_\beta/U_o = 0.629$.³ (The lower bound is not achieved because the extreme border conditions $c_1 = 0$ or $c_2 = 0$ are not included.) However, the average ratio is $U_\beta/U_o = 0.982$ and the median ratio is $U_\beta/U_o = 1.000$. We conclude that for the majority of parameters, the commission contract yields nearly the optimal profit in the fixed c_2 model. In the next section we report that this also matches the numerical analysis of the main model.

To summarize the main results from the fixed c_2 model: (i) according to Proposition 1 the optimal contract has higher consumer surplus than the fixed contract if and only if “poor service” is the best version of the fixed contract, and (ii) Proposition 2 provides a lower bound for the platform’s profit with the commission contract relative to the optimal contract.

2.6. Numerical Study

To study the performance of the five contracts in our main model, we constructed 14,700 scenarios with the goal to cover the set of feasible and plausible parameters. Table 2 summarizes the parameters used to create the scenarios. Without loss of generality, the demand intercept is set to $\bar{a} = f_l a_l + f_h a_h = 100$. The two demand states are $a_l = \delta \bar{a}$ and $a_h = (2 - \delta) \bar{a}$, which includes from a minimal level of variance in demand outcomes ($\delta = 0.9$) to nearly the maximal variance ($\delta = 0.1$). The probability of the low demand

³These 3,600 cases are constructed from the following combinations: $f_h = 0.05$; $a_l/\bar{a} = \{0.1, 0.2, \dots, 0.9\}$; $\bar{a} = 100$; $a_h = 200 - a_l$; $c_2/a_l = \{0.025, 0.075, \dots, 0.975\}$; $c_1/\hat{c}_1 = \{0.025, 0.075, \dots, 0.975\}$, where $\hat{c}_1 = f_h(a_h - a_l) + (a_l - c_2)$ is the maximum feasible value for c_1 .

state ranges from a low of 0.05 to the high of 0.95. (Proposition 2 suggests that the f_l and $f_h = 1 - f_l$ probabilities are important for comparing the optimal and commission contracts.) In all scenarios the provider's participation cost, c_2 , is Gamma distributed, with mean μ and standard deviation σ . The coefficient of variation of the participation cost ranges from a low 0.05 to a relatively high 1.5. The mean of the participation cost, μ , is selected relative to the average demand intercept value, \bar{a} , by adjusting $G(\bar{a})$ to correspond to a particular fractile of the distribution, ranging from 0.01 to 0.99. In the former case the average participation cost is high relative to consumer willingness to pay, i.e., $\bar{a} \ll \mu$, whereas in the latter case participation costs are relatively low, i.e., $\mu \ll \bar{a}$. Finally, the joining cost, c_1 , spans the range from a low value ($0.05\bar{c}_1$), to nearly its upper bound ($0.95\bar{c}_1$).

Parameters	Included values
δ	{0.1, 0.25, 0.5, 0.75, 0.9}
f_l	{0.05, 0.1, 0.25, 0.5, 0.75, 0.9, 0.95}
σ/μ	{0.05, 0.1, 0.25, 0.5, 1, 1.5}
$G(\bar{a})$	{0.01, 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 0.9, 0.95, 0.99}
c_1/\bar{c}_1	{0.05, 0.1, 0.25, 0.5, 0.75, 0.9, 0.95}

Table 2: Tested parameter values. All combinations of these values constitute 14,700 scenarios.

Table 3 reports on the frequency of different versions of the fixed contract. When the fixed contract serves both demand states (2,253 scenarios), it does so with one of two extreme versions. The poor service version is more common (73.8% of 2,253 scenarios) - capacity is set to the low demand state so that providers are fully utilized but demand is rationed. The other extreme is the "poor utilization" version - capacity is set to the high demand state, which never rations demand but leaves providers with poor utilization when low demand occurs. As expected, the low capacity (poor service) version is more prevalent when the joining cost, c_1 , is high, otherwise the high capacity (poor utilization) version tends to be selected. As is true in the fixed c_2 model, no scenarios were found which have both capacity and demand rationing.

Version	Number of scenarios	%
“Poor utilization” - capacity equals high demand, capacity rationing occurs	591	4.0%
“Poor service” - capacity equals low demand, demand rationing occurs	1,662	11.3%
Only the high demand state served	10,926	74.3%
Neither state served - unable to earn a positive profit	1,521	10.3%

Table 3: Frequency of different versions of the fixed contract.

2.6.1. Profit comparison

Table 4 reports (left side) on the profit performance of the four sub-optimal contracts relative to the optimal contract in all 14,700 scenarios. In this table, and in the remaining discussion, we use the subscripts f , w , p , β , and o to refer to the fixed, dynamic wage, dynamic price, commission and optimal contracts, respectively. On average, the fixed, dynamic wage and dynamic price contracts perform poorly relative to the optimal contract, earning only on average 75.5%, 76.2% and 79.1% of the optimal profit respectively. However, this is due to the very poor performance of a few scenarios: the median performance of those three contracts is considerably better: 96.6%, 97.1%, 98.1%. Furthermore, while the dynamic wage and the dynamic price contracts perform better than the fixed contract, their incremental performance on average is not substantial. This suggests that in this context it is insufficient to operate dynamically only on one dimension (price or wage). In contrast, while the commission contract is not optimal, its performance is nearly optimal - the average profit earned with the commission contract is 99.3% of the optimal profit and with 95% of the scenarios the commission contract earns at least 96.6% of the profit of the optimal profit. (A similar result is obtained in the fixed c_2 model.) However, there are a few scenarios in which the commission contract performs poorly - in the worst scenario the commission contract earns only 63.7% of the optimal profit. That performance is close to the analytical lower bound from the fixed c_2 model (Proposition 2) for these scenarios, $U_\beta/U_o = \frac{1}{2}(1 + \sqrt{0.05}) = 0.612$.

Table 4 also reports (right side) on the subsample of 2,253 scenarios in which the fixed contract serves demand in both states. These scenarios are considered to be less extreme

	U_f/U_o	U_w/U_o	U_p/U_o	U_β/U_o
Minimum	0.000	0.000	0.000	0.637
5%	0.000	0.000	0.000	0.966
25%	0.620	0.632	0.752	0.998
50%	0.966	0.971	0.981	1.000
75%	1.000	1.000	1.000	1.000
95%	1.000	1.000	1.000	1.000
Maximum	1.000	1.000	1.000	1.000
Average	0.757	0.762	0.791	0.993
	U_f/U_o	U_w/U_o	U_p/U_o	U_β/U_o
Minimum	0.000	0.000	0.005	0.824
5%	0.046	0.046	0.326	0.970
25%	0.460	0.475	0.797	0.997
50%	0.738	0.792	0.939	1.000
75%	0.904	0.943	0.983	1.000
95%	0.976	0.988	0.997	1.000
Maximum	0.995	0.997	0.999	1.000
Average	0.652	0.680	0.844	0.994

Table 4: Relative Profitability of Suboptimal Contracts. Profit performance of the four suboptimal contracts relative to the optimal contract in all 14,700 scenarios (left table) and in the 2,253 scenarios in which the fixed contract serves both demand states (right table). The subscripts f , w , p , β , and o to refer to the fixed, dynamic wage, dynamic price, commission and optimal contracts, respectively.

(and therefore more plausible) because the variance in demand is not so large and provider cost is not so high as to cause the platform to restrict attention exclusively to a single demand state. In this sample, three of the sub-optimal contracts perform worse than in the broader sample. Adding only dynamic wages to the fixed contract provides only a marginal improvement, whereas adding only dynamic pricing boosts the platform's profit considerably. However, there are substantial losses in profit even with dynamic pricing. In contrast, the commission contract improves its performance in this sample, in particular its worst case performance is better (yielding 82.4% of optimal profit).

It is worth emphasizing that the fixed contract performs poorly relative to the optimal contract (or the commission contract) because it charges too little during high demand and it charges *too much* during low demand. The popular press likes to emphasize higher prices during peak demand periods, but it is important to recognize that a fixed price leads to poor utilization among providers during low/normal demand and that destroys some value in the system, value that can be recaptured through the use of dynamic pricing. Thus, while consumers may (understandably) dislike the elevated prices paid during high demand, they should appreciate the benefit of paying a lower price when low/normal demand prevails.

The overall conclusions from these results are (i) it is insufficient to dynamically adjust only wage or only price, i.e., the platform should adjust both price and wage in response to demand and (ii) although the commission contract constrains the platform with the requirement of a fixed ratio between wage and price, the platform is nevertheless able to earn nearly the optimal profit in the vast majority of scenarios.

2.6.2. Membership fee contract

Although the commission contract is nearly optimal in the vast majority of cases, it is worth asking if there exists another simple contract that might perform even better. One option is a membership fee contract that has been applied in several industries (Rochet and Tirole (2006)) and has been specifically suggested for ride-sharing (Economist (2014)).

With a membership fee contract the platform sets dynamic prices, providers keep all of the revenue they earn (as in a 100% commission) and the platform earns revenue by charging providers a fixed fee to join the platform. Providers join the platform only if their earnings net of the joining fee exceeds their requirement, c_1 , and participation behavior continues to be governed by rational expectations. Unfortunately, the membership contract lacks a mechanism to limit excessive participation in the low demand state, which is an important feature of the commission and optimal contracts. Consequently, there can be a considerable loss in system value and that limits the platform’s potential earnings. Let U_m be the platform’s best profit with the membership fee contract. In our preferred sample of 2,253 scenarios, the median ratio of the platform’s profit with the membership fee contract to the optimal profit, U_m/U_o , is only 0.858 and the lowest ratio is 0.565. Thus, the membership fee contract is not a suitable alternative to the commission contract. (Details to evaluate the membership fee contract are available from the authors.)

2.6.3. Consumer, provider and system surplus

Turning to consumer surplus, we use the fixed contract as the benchmark. Tables 5 and 6 provide consumer surplus results for the set of scenarios with poor utilization or poor service with the fixed contract. The impact of adding a dynamic component to the fixed contract depends starkly on which component is made dynamic. If dynamic wages are added to the fixed contract, then consumers are always better off (i.e., $1 < S_w/S_f$ in all cases). To explain, the fixed contract with poor utilization mitigates the capacity rationing in the low demand state by constraining recruitment. Restricting recruitment limits the excess participation in the low demand state that causes capacity rationing. Once a dynamic wage is allowed, the platform can mitigate capacity rationing in the low demand state by lowering the wage in that state. This enables the platform to increase recruitment, which is beneficial to consumers. Similarly, the fixed contract with poor service substantially restricts recruitment to eliminate capacity rationing. But then a considerable amount of demand rationing occurs in the high demand state. The addition of dynamic wages allows

fractile	S_w/S_f	S_p/S_f	S_β/S_f	S_o/S_f	N_w/N_f	N_p/N_f	N_β/N_f	N_o/N_f
Minimum	1.001	0.333	0.723	0.706	0.847	0.539	0.603	0.629
5%	1.003	0.541	0.780	0.777	0.871	0.686	0.748	0.756
50%	1.025	0.854	0.957	0.956	0.995	0.911	0.945	0.946
95%	1.130	0.975	0.992	0.992	1.043	0.985	0.989	0.989
Maximum	1.234	0.986	0.994	0.994	1.099	0.989	0.993	0.994

Table 5: Relative Consumer Surplus with Poor Utilization. The ratio of consumer surplus and number of providers with the dynamic wage, dynamic price, commission or optimal contract to the fixed contract in the 591 scenarios with poor utilization.

the platform to increase the number of recruited providers while ensuring that providers continue to be fully utilized in both demand states. The increase in recruitment again benefits consumers.

Although adding dynamic wages is beneficial to consumers, the same cannot be said of dynamic prices (i.e., $S_p/S_f < 1$ in all cases). This is particularly evident with the fixed contract with poor service (Table 6). In this case, dynamic prices can address demand rationing without changing recruitment or the wage: the platform simply increases the price in the high demand state so that demand in both states matches the number of providers willing to participate under the fixed wage. The same number of consumers are served, but the high price screens consumers by their willingness to pay, improving platform profit, but lowering consumer surplus. (Better screening improves consumer surplus, but always by less than the loss of consumer surplus due to a higher price.) Dynamic prices are also problematic for consumers with the fixed contract with poor utilization (Table 5). In this case the fixed contract selects an intermediate wage and price which results in too little demand in the low demand state and too much demand in the high demand state. The addition of dynamic prices allows the platform to let its prices diverge - a low price in the low demand state and a high price in the high demand state. Increasing price in the high demand state reduces the maximum demand, so the platform can offer a smaller wage and recruit fewer providers. Neither the reduction in available supply nor the higher price benefits consumers.

fractile	S_w/S_f	S_p/S_f	S_β/S_f	S_o/S_f	N_w/N_f	N_p/N_f	N_β/N_f	N_o/N_f
Minimum	1.000	0.001	1.001	1.001	1.000	1.000	1.014	1.015
5%	1.002	0.115	1.005	1.005	1.000	1.000	1.028	1.029
50%	1.053	0.716	1.138	1.128	1.025	1.000	1.283	1.298
95%	1.580	0.952	3.962	3.912	1.336	1.000	6.335	6.378
Maximum	2.644	0.976	190.175	190.016	1.975	1.944	360.601	360.491

Table 6: Relative Consumer Surplus with Poor Service. The ratio of consumer surplus and number of providers with the dynamic wage, dynamic price, commission or optimal contract to the fixed contract in the 1,662 scenarios with poor service.

The optimal contract combines the dynamic wage contract, which is good for consumers, with the dynamic price contract, which is bad for consumers. Consequently, the optimal contract presents a mixed result for consumers, but one with a clean demarcation - consumers are better off with the optimal contract if the fixed contract chooses the poor service version (Table 6) and consumers are worse off with the optimal contract if the fixed contract chooses the poor utilization version (Table 5). Proposition 1 yields the same result for the fixed c_2 model

The commission contract provides nearly the same consumer surplus as the optimal contract, which is to be expected given that the two contracts yield similar surplus (i.e., profit) for the platform. Furthermore, as the poor service version of the fixed contract is more likely as the joining cost increases, it is expected that consumer surplus with the commission contract is more likely to increase relative to the fixed contract when the selected commission rate is high because the platform offers a high commission generally when providers incur high joining costs. Figure 7 confirms this intuition. The figure plots consumer surplus with the commission contract relative to the fixed contract (y-axis) as a function of the selected commission (x-axis). While there is variation, the general pattern is clear - as the commission rate increases, consumers are more likely to be better off with the commission contract than the fixed contract. As a point of reference (and with the understanding that our model is stylized), ride-sharing platforms tend to offer an 80% commission. Among the 864 scenarios that select a commission of 80% or higher, consumers surplus with the commission contract is higher than with the fixed contract in 859 scenarios (or 99.4% of

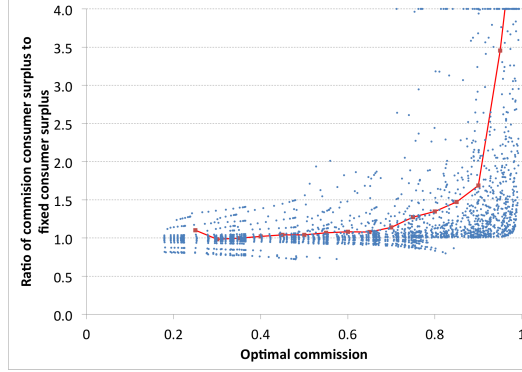


Figure 7: Relative Consumer Surplus as a Function of Commission. The ratio of consumer surplus with the commission contract to consumer surplus with the fixed contract as a function of the commission earned by providers with the commission contract in the 2,253 scenarios in which the fixed contract serves both demand states. Squares indicate the average ratio for scenarios grouped by the commission contract commission in 0.05 intervals.

them) and always higher whenever the commission is 82% or higher.

Tables 5 and 6 also report on provider surplus, which equals Nc_1 with all contracts. Thus, provider surplus is determined by the number of providers who join the platform, N . As with consumers, whether providers are better off from a switch from the fixed contract to the optimal or the commission contract depends on which of the two versions of the fixed contract is adopted. The poor utilization version of the fixed contract recruits too many providers relative to the optimal, so the optimal contract reduces the number of providers, decreasing their total surplus. In contrast, the poor service version of the fixed contract does not recruit enough providers, so total provider surplus increase with a switch to the optimal (or commission) contract.

2.7. Self-Scheduling vs. Central-Scheduling of Capacity

As an alternative to self-scheduling (providers deciding when it is best to participate), the platform could decide how many and which providers participate, a practice we call “central-scheduling.” The key advantage of central-scheduling is that it allows the platform to eliminate the inefficiency of capacity rationing: the platform would never choose to

have more providers working than necessary as that lowers the providers' earnings, making recruiting them more costly. It can also assist with demand rationing: if the number of providers on the platform exceeds demand, then all demand can be served. However, the key limitation of central-scheduling is that the platform does not observe the providers' participation costs. It would simply be too costly to credibly learn the details of every provider's planned outside activities at every possible moment. Consequently, when the platform uses central-scheduling, it can regulate the number of providers who participate, but it must select a random sample of providers, which may not be the set with the lowest participation costs. Providers anticipate that they may be scheduled to participate at less than ideal times, which affects their decision to join the platform.

The optimal contract with self-scheduling use dynamic prices and wages to eliminate capacity and demand rationing (given the pool of providers who join, N). Thus, central-scheduling is not advantageous relative to self-scheduling in terms of capacity and demand rationing, but it suffers the disadvantage of not being able to select the providers who have the lowest participation costs - because providers lack control over when they participate, they demand higher compensation to join the platform and the platform is forced to recruit fewer providers. Consequently, it is straightforward to prove that the platform earns higher profit and providers earn higher surplus with self-scheduling than with central-scheduling of capacity.

In contrast, Gurvich, Lariviere, and Moreno (2015), show that self-scheduling is less profitable for a platform than central-scheduling. Unlike Gurvich, Lariviere, and Moreno (2015), in our model providers make joining decisions based on rational expectations of their future earnings. This forces the platform to internalize the costs faced by providers. Hence, because providers value the flexibility of self-scheduling, so does the platform.

Based on our sample of 14,700 scenarios from the numerical study, the platform's best profit with central-scheduling is only 35.7% of the best profit with self-scheduling providers on average. Providers earn only 33.6% on average with central-scheduling relative to self-

scheduling. In sum, self-scheduling, by allowing providers to self-select when it is best to participate, is considerably better for the platform and providers than central-scheduling.

2.8. Discussion

Our model captures some key features of platforms with self-scheduling capacity. In particular, demand and capacity rationing can occur because demand varies considerably over time (high and low demand periods), long run capacity is rigid and too many providers may choose to participate, thereby destroying rents in the short term and reducing the attractiveness of joining the platform in the long term. However, our model abstracts away from a number of other issues that affect these platforms in practice. We discuss several possible extensions in this section that merit further investigation.

We assume there exists a large pool of potential providers who all require at least c_1 in expected profit as a threshold before they are willing to join. Once providers join, we assume they all can provide the same amount of capacity to the platform in period 2. In practice, there is heterogeneity in the wages a provider requires to join the platform and heterogeneity in the number of hours they are willing to work. It is possible to add heterogeneity in c_1 to our model in the form of a two point distribution: there are M providers with joining cost c_l and an unlimited number with a higher joining cost, c_h . For M sufficiently large, the best version of all contracts remains the same as if $c_1 = c_l$. For M sufficiently small, the marginal provider has a joining cost of c_h and earns zero surplus from joining the platform, while the M providers with the lower joining cost, c_l , enjoy some surplus from joining. Due to this increasing supply curve, we anticipate that the profit and surplus gaps between the fixed contract and the optimal contract are reduced relative to our observations with a fixed c_1 : the optimal (and commission) contract benefits from increased recruitment of providers, but an increasing supply curve mitigates the optimal contract's ability to take advantage of recruiting a larger pool of providers.

In our model the platform does not incur explicit recruiting costs and providers do not quit

the platform once they join. Furthermore, there is no learning in our model - providers correctly anticipate their future earnings. In practice, platforms are indeed concerned with provider recruitment costs and retention. Such issues could influence how the platform matches customers to providers - we assume random matching but that may not be the best for a platform that wants to manage retention.

Our platform neither faces competition from another platform or from other firms offering similar services. Even with competition it is important for the platform to recruit the correct number of providers and to ensure that they are utilized properly. But competition could alter the attractiveness of the contracts we consider, both in terms of the competition for customers as well as the competition for providers. For example, Liu and Zhang (2012) show that competing firms may prefer to commit to fixed pricing rather than dynamic pricing.

We use a single joining period to represent long-term capacity decisions and a single period to represent short-term participation decisions. These are most appropriate when a platform has achieved steady-state and providers make many participation decisions which are both similar and uncorrelated. In practice a platform may experience growth over time, which should be represented with multiple joining periods. Similarly, one could consider a model with multiple short-term participation decisions. Such a model would allow the investigation of the impact of demand correlation over time as well as correlation between demand and participation costs (e.g., “high demand” in a period could be associated with “high participation” costs for providers).

2.9. Conclusion

We study a platform that offers a service via a pool of independent providers. Providers self-schedule when they offer their service to the customers on the platform and decide whether or not to join the platform based on their earnings expectations. Demand varies over the long-term but is predictable in the short-term. Two inefficiencies can arise: (i) demand

can be rationed either because too few providers join the platform or too few choose to participate; and (ii) capacity can be rationed because competition for a limited number of jobs leads too many providers to participate. Demand rationing is costly because some customers are unable to access the service that they value at the price charged, and the customers that do get the service might not be the ones that value it the most. Capacity rationing is costly because participating providers are not fully utilized but still incur their full opportunity cost of joining the platform. Both forms of rationing factor into the decision of providers as to whether to join the platform or not.

Although self-scheduling removes some control from the platform (it cannot directly control the number of providers who work), it allows providers to self-select when it is most beneficial for them to work. We show that this additional flexibility is beneficial to providers, the platform and consumers.

We study several contractual forms that vary in whether prices and/or wages respond to demand. The most basic contract, the fixed contract, sets a single price and wage no matter what demand level occurs. To the fixed contract the platform could add either dynamic wages or dynamic prices. The optimal contract requires that the platform chooses both a price and a wage contingent on demand. We find that adding one dynamic component to the fixed contract (either wage or price but not both) increases the platform's profit but still leaves the platform with substantially lower profit than what it could earn with the optimal contract, which is dynamic in both components. A commission contract chooses both price and wage dynamically, but includes the added constraint of a fixed ratio between the two. The commission contract mimics pricing used in practice, such as Uber's surge pricing. Our main result is that even though the commission contract is not optimal, it yields nearly the optimal profit for the platform in the vast majority of plausible scenarios.

While maximizing profit is clearly an important objective for the platform, it isn't the only relevant one. A considerable amount of controversy has arisen over whether self-scheduling providers should be treated like employees (e.g. given additional rights and benefits) and

whether surge pricing gouges consumers. Hence, a platform should also be concerned with how it influences both provider and consumer surplus.

The optimal contract leads to ambiguous welfare implications, which depend on how the fixed contract manages demand and capacity. If providers are relatively inexpensive (i.e. their opportunity cost to join the platform is low), then the fixed contract recruits an ample number of providers and underutilizes them during low demand periods. Adding dynamic prices and wages to that situation always works to the disadvantage of providers and consumers because the platform recruits fewer providers and, in the high demand state, charges more and serves fewer customers. However, if providers have a high opportunity cost, then the fixed contract recruits a limited number of providers and forces customers during peak demand to suffer through poor service. In those cases, providers and consumers are better off with the introduction of dynamic prices and wages: capacity expands to serve more customers in all demand states. To frame this in the context of ride-sharing, if with the fixed contract (e.g. taxis) it is hard to find service at peak demand times (e.g. a rainy evening), then Uber's introduction of surge pricing (i.e., dynamic pricing and wages) is likely to make all stakeholders (Uber, drivers, and consumers) better off.

CHAPTER 3 : Distinguishing the Gig-Economy from Two-Sided Markets

3.1. Introduction

“Gig-economy” platforms like Uber, TaskRabbit, and Postmates have created marketplaces in which services are performed on-demand for consumers by independent providers. Providers decide whether, when, and how much service to offer via the platform, and customers demand immediate service when their need arises. This setting has motivated an emerging stream of literature studying the platform’s optimal strategy to coordinate its decentralized labor force.

Naturally a provider’s interest in working depends on the amount of money he earns from a completed service (Farber (2015), Chen and Sheldon (2015)). A platform may then manage its workforce indirectly by manipulating the “wage” it offers providers per service. Typically the platform pays providers a commission, meaning providers earn a percentage of the price paid by the consumer. Consequently the firm may incentivize working by manipulating either the price charged to consumers or the commission percentage. For example, Uber manipulates price through its “Surge Pricing” policy, which increases price during peak demand to entice drivers to offer rides, while Lyft manipulates the commission percentage through its “Power Driver Bonus” program, which pays a higher percentage to drivers that complete a high volume of rides during peak hours.

Providers also respond to the availability of demand to serve. In on-demand service marketplaces, a provider must already be available in order to receive customers service requests. For example, Uber drivers are matched with nearby passengers, so drivers must already be on the road to receive a ride request. As a result, providers incur the opportunity cost of their time whether or not they are serving a customer. The more time a provider expects to spend serving customers, the more likely participating in the market will be worth the opportunity cost of his time.

In broad strokes, a gig-economy marketplace resembles a two-sided market. The volume of agents on one side of the market depends on the volume of agents on the other side (e.g. the number of providers working depends on the volume of demand for service). The platform manages the volume of transactions from which it profits by manipulating the price (or wage) charged to each side. Furthermore, agents may be heterogeneous in their valuation of the transaction (e.g. customers have heterogeneous value of receiving service) and in their valuation of membership in the platform (e.g. providers have heterogeneous opportunity cost of time).

However, the classic literature studying two-sided market limits its analysis in some important ways: it does not consider agent capacity constraints or inter-agent competition (e.g. Rochet and Tirole (2003), Armstrong (2006), Rochet and Tirole (2006), Weyl (2010)). These limitations are appropriate in applications where participation in the platform's market is a long term decision (e.g. markets for credit cards, video game consoles). In these settings, consumers seek many interactions via the platform and are largely variety seeking. For example, a credit card holder wants use his card at the grocery store, at the barber, and at the veterinarian, while the gamer wants to play many different games on her console. Consequently, service providers expect to interact with many consumers, and their interaction with a consumer does not cannibalize interactions between that consumer and other providers. Further, because consumers remain in the market for a long time, providers have the opportunity to completely satisfy demand even in the face of short term capacity constraints. For example, a video game producer may backorder a consumer's request for their product, but it is likely that the consumer will still be using the same gaming console by the time she receives her delayed order.

In contrast, gig-economy customers and providers make decisions about participation on a platform on a much shorter time scale. Consequently, the capacity constraints faced on both sides of the market become relevant: customers can demand only one service at a time and providers can offer service to only one customer at a time. Furthermore, gig-

economy providers are largely homogeneous from the consumer’s point of view, so providers compete for limited demand. Extant work in this area captures this by characterizing the volume of transactions from which the platform profits as the number of successful provider-customer matches, and the volume of providers as a function of their expected utilization. Banerjee et al. (2015) propose a queueing model of on-demand services where the volume of providers depends on the expected time until a successful match. Chen and Hu (2016) study a dynamic model in which each provider decides how long to wait before offering service, a decision based the probability of successfully matching at each time t . Otherwise, the extant literature captures the dependence of provider volume on aggregate demand via the ratio of demand volume to supply volume (e.g. Cachon et al. (2017), Hu and Zhou (2017), Bai et al. (2016), Taylor (2016)). Keeping with this precedent, the model below links provider volume to the ratio of demand and supply.

The goal of much of the gig-economy literature is to understand pricing in a gig-economy setting. Of particular interest is dynamic pricing. Dynamic prices may be used to solve mismatches in supply and demand arising from stochastic customer and provider arrival processes. However, Banerjee et al. (2015) and Chen and Hu (2016) demonstrate that gains from dynamic prices in this setting are small. Dynamic prices may also be used to correct mismatches in supply and demand across states of the world (e.g. peak and off-peak times). As shown by Cachon et al. (2017), dynamic prices dramatically improve platform profit in this setting. It is then natural to ask whether the common practice of paying providers a fixed percentage of this dynamic price, called a “fixed commission,” is a good policy relative to a dynamic commission. Bai et al. (2016) provides examples of the profit loss resulting from a fixed commission, Bimpikis et al. (2016) studies the consequences of spatial layout on the cost imposed by a fixed commission, and Cachon et al. (2017) and Hu and Zhou (2017) provide bounds on the relative profitability of a platform with a fixed commission and a platform with a dynamic commission. Bimpikis et al. (2016), Cachon et al. (2017), and Hu and Zhou (2017) both demonstrate that a fixed commission can be optimal in some cases, while Cachon et al. (2017) and Hu and Zhou (2017) show that in many cases it reasonably

approximates platform profit with a dynamic commission.

The purpose of this paper is to determine whether the modeling differences between the classic two-sided market literature and the emerging gig-economy are significant. Since much extant gig-economy work focuses on the performance of the fixed commission contract, my outcome of interest is the profitability of a fixed commission relative to a state-dependent commission. I consider a platform that faces two states of the world, which are distinguished by their latent demand (i.e. the number of customers demanding a free service) and latent supply (i.e. the number of providers offering service for an infinite wage). I restrict my attention to markets where the peak state has (i) higher latent demand than the off-peak state, (ii) insufficient latent supply to satisfy latent demand, and (iii) relatively less latent supply than the off-peak state. This captures a common dynamic faced by gig-economy platforms. For example, Uber notoriously faces dramatically higher peak demand, during which passengers must (i) wait for a ride and (ii) wait longer for a ride, indicating capacity is both scarce and scarcer than during off-peak hours Hall et al. (2015). I use latent demand and supply to measure three sources of imbalance: demand imbalance across states, supply scarcity imbalance across states, and the imbalance between supply and demand within the peak state. It is natural to expect the fixed commission to perform well both in settings with little imbalance and also in settings with extreme imbalance, which cause the platform to effectively serve just one state. The analysis that follows will demonstrate the extent to which this intuition is correct.

In this context, I construct a model of a gig-economy platform and a corresponding model of a two-sided market platform. I derive for each model the best price and commission structure with and without the fixed commission restriction. I show that the inclusion of inter-agent competition and capacity constraints cause the gig-economy platform to match supply and demand, whereas a two-sided market platform has no such objective. The profitability of matching supply and demand causes the fixed commission to respond differently to sources of imbalance with the gig-economy model than with the two-sided market model.

In particular, fixed commission performance depends on both between-state imbalances and within-state imbalance with the gig-economy model, while only depending on between-state imbalances with the two-sided market model. Numerically I show that, unexpectedly, the gig-economy fixed commission performs better with large demand imbalance across states and large peak supply scarcity than with small demand imbalance and small peak supply scarcity. This stands in contrast to the two-sided market fixed commission, which is indifferent to large and small demand imbalance and is independent of peak supply scarcity. Finally I show that, while the conditions that cause poor fixed commission performance differ across models, the fixed commission generally performs nearly optimally. Hence, it is not the efficacy of the fixed commission that distinguishes the gig-economy but the conditions that reverse this result.

3.2. Model

In platform economics, the platform coordinates the actions of independent agents. In the setting of interest, agents fall into two categories: on one side are customers demanding service, on the other are service providers. Service providers are equally able to serve each customer and customers do not have preferences over service providers. Let J denote the set of states of the world, indexed by subscript j . The platform's challenge is to manage the number of service providers and customers in each state of the world. The platform does this by charging a price, p_j , to customers and by offering a commission, β_j , to service providers. Service providers earn $\beta_j p_j$ per service.

Customers have heterogeneous valuations for the service. Customers enjoy this value and pay the platform's price each time they transact with a service provider. Denote by ϕ_j^D the number of transactions a customer expects to have with service providers. Service providers have heterogeneous costs of making themselves available to serve customers. Service providers earn a percentage of the price paid per interaction with a customer but incur the same cost regardless of the number of interactions with customers. For example, Uber drivers are only matched to customers nearby, so Uber drivers must already be on the road

to be matched with a passenger. Consequently, drivers incur their opportunity cost of time independently of the number of passengers they serve. Denote by ϕ_j^S the number of transactions a service provider expects to have with customers. Then a customer with valuation v demands service in state j if

$$(v_j - p_j)\phi_j^D \geq 0$$

and a service provider with cost c offers service in state j if

$$\beta_j p_j \phi_j^S - c \geq 0.$$

The platform earns its margin on each transaction between a customer and a service provider. The platform wishes to maximize its profit, π , which is the product of its margin, $(1 - \beta_j)p_j$, and the volume of transactions, V_j :

$$\pi \doteq \sum_{j \in J} (1 - \beta_j)p_j V_j.$$

The purpose of this analysis is to compare the platform's performance using the definition of transaction volumes in a gig-economy setting versus in a traditional two-sided market setting. With a classic two-sided market model, the volume of transactions a customer (service provider) has is linearly increasing in the number of active service providers (customers), and the total number of transactions from which the platform profits is the product of the populations on either side of the market. In contrast, a gig-economy model captures a service provider's capacity constraint (normalized to a single customer), and a customer's desire for a single service. Agents in the gig economy respond the volume of other-side agents through its effect on the probability of that the agent is successfully matched.

Of particular interest to the emerging literature studying gig-economy operations is the performance of a *fixed commission* (e.g. Bai et al. (2016), Banerjee et al. (2015), Cachon et al. (2017), Hu and Zhou (2017)). With a fixed commission, the platform restricts its

wage in state j to be constant multiple, β , of price. This interest is motivated by the prevalence of this pricing scheme in practice (e.g. Uber, Airbnb, Postmates). The analysis below echoes this interest. In each setting (e.g. gig-economy, two-sided market), I analyze the model with and without the fixed commission restriction.

The analysis below is based on the following standard assumptions. Let d_j denote latent demand, the mass of customers interested in service were service free in state j . Assume that customer valuations are independent and identically distributed with cumulative distribution function F and probability density function f . Let s_j denote latent supply, the mass of service providers willing to serve if the wage were infinite in state j . Assume that service provider costs are independent and identically distributed with cumulative distribution function G and probability density function g . Then demand for service in state j is $D_j = d_j \bar{F}(p_j)$ and supply of service in state j is $S_j = s_j G(\beta_j p_j \phi_j)$. I further assume $F(0) = G(0) = 0$, and f and g are log-concave, which characterizes many common distributions (for more details, see Bagnoli and Bergstrom (2005)).

Let us assume that the platform faces only two states of the world, called “low” and “high.” In the high state, the platform faces higher demand than in the low state and has insufficient capacity to satisfy demand, i.e. $d_l < d_h$ and $s_h < d_h$. Furthermore, capacity is relatively less available in the high state, i.e. $s_h/d_h < s_l/d_l$. This captures a common dynamic faced by gig-economy platforms. For example, Uber notoriously faces dramatically higher peak demand, during which passengers must (i) wait for a ride and (ii) wait longer for a ride, indicating that capacity is both scarce and scarcer than during off-peak hours (Hall et al., 2015). The following parameters, all of which belong to the interval $[0, 1]$, measure a source of imbalance for the platform:

$$\begin{aligned}\alpha &\doteq d_l/d_h \\ \rho &\doteq s_h/d_h \\ \gamma &\doteq \frac{s_h/d_h}{s_l/d_l}.\end{aligned}$$

The first measures imbalance in latent demand across states; the second measures imbalance in latent supply and demand in the high state; the third measures the imbalance in latent supply scarcity across states. It is natural to expect the fixed commission to perform well both in settings with little imbalance and also in settings with extreme imbalance, which cause the platform to effectively serve a single state. In the analysis that follows, I identify the role that each of these parameters play in the performance of the fixed commission.

3.3. Gig Economy

In the gig-economy, customers demand a single service and service providers cannot serve more than one customer. The platform randomly matches available service providers to customers requesting service. Customers and service providers must decide whether to request service and make themselves available to serve, respectively, in advance of matching. Consequently, $\phi_j^D = \min\{S_j/D_j, 1\}$ is the probability that a customer receives service (conditional on being willing to pay p_j). Similarly, $\phi_j^S = \min\{D_j/S_j, 1\}$ is the probability that a service providers receives a customer to serve (conditional on being available). As described above, $D_j = d_j \bar{F}(p_j)$ and S_j is implicitly defined by the relationship $S_j = s_j G(\beta_j p_j \min\{D_j/S_j, 1\})$. The platform profits from the total volume of matches, $V_j = \min\{D_j, S_j\}$.

First consider the solution to the platform's optimal price and commission decision. As described in Proposition 3, at optimal the platform uses price to match supply and demand in every state of the world. The platform's resulting objective is quasiconcave in β_j with a guaranteed interior solution.

Proposition 3. *The platform's optimal price and commission structure is uniquely defined by p_j^o and β_j^o satisfying the following conditions:*

$$\begin{aligned} d_j \bar{F}(p_j^o) &= s_j G(\beta_j^o p_j^o) \\ (1 - \beta_j^o) p_j^o &= G(\beta_j^o p_j^o) / g(\beta_j^o p_j^o) + \bar{F}(p_j^o) / f(p_j^o) \end{aligned}$$

Notice that the platform's choice of commission depends on the state only through p_j^o .

Consequently, if p_j^o does not vary across states, then the commissions in all states are equal.

Corollary 1. *The fixed commission is optimal if $s_j/d_j = s_{-j}/d_{-j} \forall j$.*

Stated differently, the platform chooses a fixed commission when $\gamma = 1$. If there is no variation across states in latent supply scarcity, then the platform has no use for a dynamic commission. This is a direct result of the platform's desire to match supply and demand in each state.

It is, however, common for a platform to experience such imbalance in supply scarcity across states. If the platform were restricted to offer the same commission in every state even with $\gamma < 1$, then the platform's price and commission structure is described by the following:

Proposition 4. *There exists an optimal commission, β^c in the interval $[\min_j\{\beta_j^o\}, \min_j\{\beta_h^o\}]$.*

For any β^c , the platform's chooses price to be p_j^c , which is uniquely defined by

$$\max\{\hat{p}(s_j/d_j, \beta^c), p^*\}$$

where $\hat{p}(s_j/d_j, \beta) : D_j = S_j$ and $p^ = \bar{F}(p^*)/f(p^*)$.*

The best fixed commission ensures that demand is satisfied in each state of the world while offering a commission no smaller than the smallest dynamic commission but no larger than the largest dynamic commission. While there exists a special case in which a platform experiences no profit loss from fixing its commission, in general the fixed commission restriction costs the platform something. Existing literature demonstrates that in many cases this cost is not too great (e.g. Cachon et al. (2017), Hu and Zhou (2017)). However, the literature also demonstrates that a platform can lose nearly half of its possible profit by choosing to fix its commission (Cachon et al., 2017). In the analysis that follows, I will illustrate which market conditions make the fixed commission costly. In particular, I will study the effect of different measures of imbalance, α , γ , and ρ .

Define platform profit with and without the fixed commission restriction to be π_g^c and π_g^o

respectively. Then measure the gap in profit by

$$P_g \doteq \pi_g^c / \pi_g^o.$$

Theorem 12. *P_g has the following properties:*

1. P_g is maximized at extreme an value of α , i.e. $\alpha = 0, 1$
2. $P_g \geq \max_j \{\pi_{g,j}^o / \pi_g^o\}$, which is maximized at an extreme value of γ and ρ

The first result parallels intuition. On the one hand, removing demand imbalance intuitively improves fixed commission performance. On the other hand, with $\alpha = 0$, the platform faces just one state of the world and the fixed commission must be optimal.

The second result provides a bound on fixed commission performance. This bound indicates that the fixed commission achieves at least half of the platform's optimal profit. Note that this bound does not require restrictions like linear demand curve (Cachon et al., 2017) or concave supply curve (Hu and Zhou, 2017).

Of further interest is the bound's dependence on γ and ρ . Due to the non-monotonicity of β_j^o in γ and ρ , evaluation of P_g 's dependence on those parameters is intractable. Instead, I turn to the bound in Theorem 12 to illustrate the effect of γ and ρ on fixed commission performance. Theorem 12 shows that the bound is maximized at an extreme value of γ and ρ . Again, this parallels the intuition that small imbalances benefit the fixed commission and large imbalances cause the platform to effectively serve just one state, making the fixed commission optimal.

3.4. Two-Sided Markets

In the traditional two-sided market literature, an agent's utility is linearly increasing in the number of other side agents. This model is tailored to settings in which membership in the market lasts significantly longer than in a gig-economy market. To see how well this model

extends to the gig economy, consider the following model. A customer in state j with value v has utility

$$(v - p_j)S_j$$

and a service provider with cost c in state j has utility

$$\beta_j p_j D_j - c.$$

This is a special case of the classic model in Rochet and Tirole (2006), where customers have no membership benefit and providers have no per-transaction benefit beyond their associated monetary gain. The resulting populations on either side of the market are $D_j = d_j \bar{F}(p_j)$ and $S_j = s_j G(\beta_j p_j D_j)$. The volume of transactions from which the platform profits is then $D_j S_j$, so the platform solves

$$\max_{p,w} \sum_j (1 - \beta_j) p_j D_j S_j.$$

Proposition 5. *As before, define $p^* = \bar{F}(p^*)/f(p^*)$. The platform's optimal price and commission structure is uniquely defined by p_j^o and β_j^o satisfying the following conditions:*

$$(1 - \beta_j^o) p_j^o D_j = G(\beta_j^o p_j^o D_j) / g(\beta_j^o p_j^o D_j)$$

$$p_j^o = p^*$$

Notice that the optimal price is independent of both the state and the platform's choice of commission. This contrasts with the importance of dynamic prices for profit maximization in the gig-economy model, as demonstrated by Cachon et al. (2017). The state-independence of optimal price also means the platform offers the same commission in all states when differences between states are the result of supply side effects.

Corollary 2. *A fixed commission is optimal with a classic model if $d_j = d_{-j} \forall j$.*

Stated differently, the platform elects to offer a fixed commission when $\alpha = 1$. When states

are differentiated only based on their latent supply, then the platform has no use for a dynamic commission. However, platforms typically face variability in latent demand. If the platform were restricted to offer the same commission in every state while $\alpha < 1$ then the platform's price and commission structure is described by the following:

Proposition 6. *There exists an optimal commission, β^c in the interval $[\min_j\{\beta_j^o\}, \max_j\{\beta_j^o\}]$. The corresponding best price in state j is $p_j^c = p^*$.*

The fixed commission restriction does not change the platform's pricing structure. The resulting fixed commission balances profit loss in each state and consequently must lie between the smallest and the largest optimal commission.

Corollaries 1 and 2 indicate that it is possible for a fixed commission to be optimal with a gig-economy model but not with a two-sided market model, and vice versa. To better understand the profit gap created by the fixed commission restriction in a two-sided market, define the profit gap by

$$P_t \doteq \pi_t^c / \pi_t^o$$

where π_t^c and π_t^o refer to profit with the optimal fixed and dynamic commissions, respectively.

Theorem 13. *P_t has the following properties:*

1. P_t is maximized by an extreme value of α , i.e. $\alpha = 0, 1$
2. P_t is maximized by an extreme value of γ
3. P_t is independent of ρ

Theorem 13 illustrates similarities and differences between the performance of the fixed commission with the gig-economy and the traditional two-sided market model. Like the gig-economy, the fixed commission performs best at extreme values of α . As before, the platform effectively faces a single state of the world when $\alpha = 0$. From Corollary 2, the fixed commission is optimal when $\alpha = 1$. Also as in the gig-economy, the fixed commission per-

formance is maximized at an extreme value of γ . Unlike the gig-economy, the performance of the fixed commission does not depend ρ . Because the gig-economy platform wishes to match supply and demand in all states, fixed commission performance depends not only on the variation in supply scarcity across states but also on the magnitude of the imbalance in supply and demand within each state. The two-sided market platform, however, does not seek to match supply and demand. Consequently, the fixed commission’s performance is independent of within state imbalance, i.e. P_t is independent of ρ . This represents a key difference in the behavior of the fixed commission with a gig-economy model versus and two-sided market model.

3.5. Numerical Analysis

To explicitly compare the performance of the fixed commission with a gig-economy model and a two-sided market model, I construct 104,976 scenarios designed to cover the set of feasible parameters. Table 7 summarizes the parameter values used in these scenarios. I normalize d_h to 1 and vary the values of d_l , s_h , and s_l by varying the values of α , γ , and ρ , all of which must belong to the interval $[0, 1]$. I assume that both consumer demand and provider cost are distributed according the gamma distribution. Each distribution is defined by a mean μ_k and a coefficient of variation CV_k , where the subscript denotes the distribution to which the quantity belongs. For each distribution, the coefficient of variation ranges from 0 to 1.

Parameters	Values
γ	{.1, .2, .3, .4, .5, .6, .7, .8, .9, .99}
α	{.1, .2, .3, .4, .5, .6, .7, .8, .9, .99}
ρ	{.1, .2, .3, .4, .5, .6, .7, .8, .9, .99}
μ_f	{.5, 1, 5}
CV_f	{.25, .5, .75, .9}
μ_g	{.5, 1, 5}
CV_g	{.25, .5, .75, .9}

Table 7: A summary of tested parameter values. All combinations of these values constitute 104,976 numerical experiments.

First observe that the gig-economy fixed commission captures most of the profit generated

by the dynamic commission. Specifically the fixed commission yields 99.58% of optimal profit on average. This replicates the findings in Cachon et al. (2017) which demonstrate that, in the gig-economy, the fixed commission is generally a good approximation of platform profit without the fixed commission restriction.

The efficacy of the fixed commission, however, is not a phenomenon unique to the gig-economy. The second row in Table 8 demonstrates that the fixed commission also captures most of the platform’s optimal profit with the two-sided market model, yielding 99.60% of optimal profit on average.

Both models experience poor fixed commission performance in a minority of cases. The third row in Table 8 indicates that the conditions that create poor performance with one model are not the same with the other. Therefore to understand the differences between these two models, it is necessary to understand the conditions that create poor fixed commission performance with each model.

	Minimum	Q1	Median	Mean	Q3	Maximum
P_g	.8876	.9973	.9997	.9958	1.000	1.000
P_t	.7289	.9993	.9999	.9960	1.000	1.000
P_g/P_t	.8877	.9985	.9999	1.0001	1.0001	1.3690

Table 8: Quartiles and mean of P_g , P_t and P_g/P_t .

To understand these settings, consider the plots in Figure 8. Each illustrates the spectrum of possible behaviors of fixed commission performance as a function of each measure of imbalance. Each plot is maximized at an extreme value of the imbalance measure, as predicted by Theorems 12 and 13. The numerical experiments are useful for identifying which extreme maximizes fixed commission performance with each model.

Considering the first column, the two-sided market fixed commission is maximized at both extremes of α . This is consistent with the intuition that little imbalance (i.e. $\alpha = 1$) produces good fixed commission performance, and that large imbalance (i.e. $\alpha = 0$) causes the platform to focus on only one state of the world, producing good fixed commission perfor-

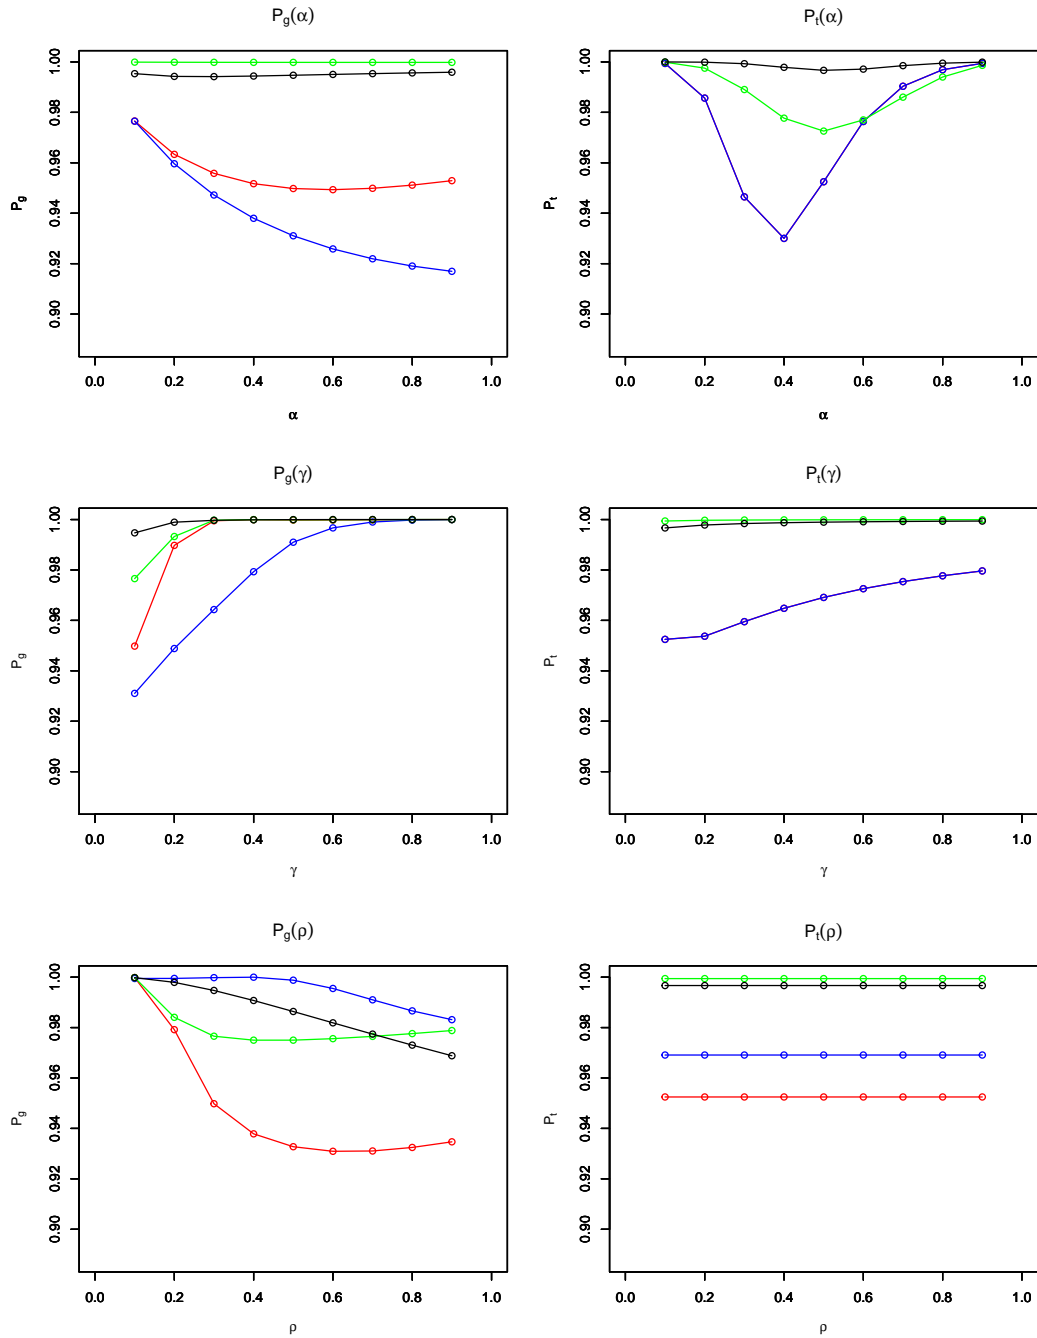


Figure 8: Illustrative plots. All plots have $\mu_g = 1$, $CV_g = .5$, and $CV_f = .25$. In the first row, curves are defined by $\{\gamma, \rho, \mu_f\}$, where red has $\{.1, .3, 5\}$; blue has $\{.1, .7, 5\}$; green has $\{.6, .3, 5\}$; black has $\{.1, .3, 1\}$. In the second row, curved are defined by $\{\alpha, \rho, \mu_f\}$, where red has $\{.5, .3, 5\}$; blue has $\{.5, .7, 5\}$; green has $\{.1, .3, 5\}$; black has $\{.5, .3, 1\}$. In the third row, curves are defined by $\{\alpha, \gamma, \mu_f\}$ where red has $\{.5, .1, 5\}$; blue has $\{.5, .5, 5\}$; green has $\{.1, .1, 5\}$; black has $\{.5, .1, 1\}$.

mance. The gig-economy fixed commission is maximized at only the latter extreme. This results from the supply-demand matching that happens in the gig-economy model. Removing latent demand imbalance across states is insufficient to allow the platform to match supply and demand with a single price in all states. Said differently, if $\gamma < 1$, $\alpha = 1$ is not enough to guarantee near-optimal performance of the gig-economy fixed commission. Consequently, the gig-economy fixed commission performs worse with little demand imbalance than with large demand imbalance.

Similarly, the gig-economy fixed commission is maximized with large peak state supply-demand imbalance. Though $\rho = 1$ removes this imbalance, with $\gamma < 1$ it is insufficient to produce near-optimal fixed commission performance. However, small ρ produces small *absolute* difference in the prices required to match supply and demand in each state, leading to near-optimal fixed commission performance. Hence the gig-economy fixed commission performs worse with little imbalance than with large imbalance. This dependence stands in contrast to the two-sided market fixed commission, whose performance is independent of ρ .

It is only in γ that fixed commission performance behaves similarly across models. Both models prefer less imbalance across states in supply scarcity.

Figure 8 demonstrates the following key takeaways. First, two-sided market fixed commission performance depends only on across-state imbalance. In contrast, gig-economy fixed commission performance depends both on across-state and within-state imbalance because of the supply-demand matching required to maximize platform profit in this model. Unexpectedly, gig-economy fixed commission performance is *worse* with small imbalance in α and ρ relative to large imbalance. This is again a result of the gig-economy platform's need to match supply and demand; removing demand imbalance or supply-demand imbalance in the peak state is insufficient to match supply and demand across states. Consequently, one model cannot predict poor fixed commission performance in the other. This justifies the study of a model specifically tailored to the gig-economy to understand gig-economy behavior.

3.6. Conclusion

In this paper, I build fixed and dynamic commission models in two settings. One uses traditional two sided markets assumptions (i.e. agent utility grows linearly in other side agent population, and the platform profits from all possible combinations of agents), while the other uses assumptions tailored to the gig-economy (i.e. agents are capacity constrained and compete with each other). Analysis of these models indicates that the difference in profitability of a fixed commission in one model versus the other is tied to imbalances in latent supply and demand across and within states. Fixed commission performance in the two-sided market model depends on between-state variability (and is maximized in the absence of such variability) and is independent of the mismatch in peak supply and demand. In contrast, fixed commission performance in the gig-economy model depends both on between-state differences and within-state differences and, in some cases, *improves* with greater mismatches. I numerically show that while both models generally predict near-optimal fixed commission performance, one model cannot predict poor fixed commission performance in the other.

This analysis solidifies the position of gig-economy research as distinct from the body of two-sided market literature. The capacity constraints and inter-agent competition inherent in the gig-economy create a dependence on within-state mismatches that is absent from the two-sided market model. Similarly, the profitability in the gig-economy of matching supply with demand in each state of the world allows fixed commission performance to *suffer* from small imbalance in latent demand and peak supply scarcity. These differences cause the two-sided market model to fail to predict the poor performance in the worst case scenarios of the gig-economy model. It is therefore all the more important to study the optimal form of a fixed commission offered by the likes of Uber using a model tailored to the gig-economy.

APPENDIX

A.1. Proofs of Chapter 1 Theorems

Proof of Theorem 1

Part 1. $\Pi_V^L \geq \Pi_A^L$ if and only if

$$\begin{aligned} \frac{a^2 N E \left[\frac{1}{C} \right]^2}{4(2E \left[\frac{1}{C} \right] + bN \text{Var} \left(\frac{1}{C} \right) + bNE \left[\frac{1}{C} \right]^2 - 1/C_H)} &\geq \frac{a^2 N}{4(E[C] + bN)} \\ E \left[\frac{1}{C} \right]^2 E[C] + E \left[\frac{1}{C} \right]^2 bN &\geq 2E \left[\frac{1}{C} \right] + bN \text{Var} \left(\frac{1}{C} \right) + bNE \left[\frac{1}{C} \right]^2 - \frac{1}{C_H} \\ \frac{1}{bN} (E \left[\frac{1}{C} \right]^2 E[C] - 2E \left[\frac{1}{C} \right] + \frac{1}{C_H}) &\geq \text{Var} \left(\frac{1}{C} \right) \end{aligned}$$

From Seaman and Odell (1985), we may bound $\text{Var} \left(\frac{1}{C} \right) \leq (1/C_L - 1/C_H)/4$, so the difference between the voluntary and automated profits maybe expressed by

$$\frac{1}{bN} (E \left[\frac{1}{C} \right]^2 E[C] - 2E \left[\frac{1}{C} \right] + \frac{1}{C_H}) - \frac{\frac{1}{C_L} - \frac{1}{C_H}}{4}$$

and is greater than zero when

$$E \left[\frac{1}{C} \right] \notin [\underline{y}, \bar{y}] = \left[\frac{1}{E[C]} - \frac{1}{E[C]} \sqrt{1 - E[C]K}, \frac{1}{E[C]} + \frac{1}{E[C]} \sqrt{1 - E[C]K} \right]$$

where $K = \frac{1}{C_H} - b(\frac{1}{C_L} - \frac{1}{C_H})^2/4$. Thus if $E \left[\frac{1}{C} \right] \geq \bar{y}$ then $\Pi_V^L \geq \Pi_A^L$.

Part 2. $\Pi_A^L \geq \Pi_V^L$ if and only if

$$\begin{aligned} \frac{a^2 N}{4(E[C] + bN)} &\geq \frac{a^2 E \left[\frac{1}{C} \right]^2}{4(2E \left[\frac{1}{C} \right] + bNE \left[\frac{1}{C} \right]^2 - 1/C_H)} \\ \iff 2E \left[\frac{1}{C} \right] + bNE \left[\frac{1}{C} \right]^2 - \frac{1}{C_H} &\geq E \left[\frac{1}{C} \right]^2 (E[C] + bN) \end{aligned}$$

By Jensen's Inequality we know $E[\frac{1}{C^2}] \geq E[\frac{1}{C}]^2$. It follows that $\Pi_A^L \geq \Pi_V^L$ if

$$\begin{aligned}
2E\left[\frac{1}{C}\right] + bNE\left[\frac{1}{C}\right]^2 - \frac{1}{C_H} &\geq E\left[\frac{1}{C}\right]^2 (E[C] + bN) \\
\iff 2E\left[\frac{1}{C}\right] - \frac{1}{C_H} &\geq E\left[\frac{1}{C}\right]^2 E[C] \\
\iff 0 &\geq E\left[\frac{1}{C}\right]^2 E[C] - 2E\left[\frac{1}{C}\right] + \frac{1}{C_H}. \tag{*}
\end{aligned}$$

The final expression is a concave quadratic in $E\left[\frac{1}{C}\right]$, so we can easily find the range of $E\left[\frac{1}{C}\right]$ satisfying (*):

$$E\left[\frac{1}{C}\right] \in [z, \bar{z}] = \left[\frac{1}{E[C]} - \frac{\sqrt{1 - \frac{E[C]}{C_H}}}{E[C]}, \frac{1}{E[C]} + \frac{\sqrt{1 - \frac{E[C]}{C_H}}}{E[C]} \right]$$

Clearly $E[C] \leq C_H$ so such a set of $E\left[\frac{1}{C}\right]$ always exists and has a subset in the feasible range, i.e. $\bar{z} \geq \frac{1}{E[C]}$. Furthermore, $z \leq \frac{1}{E[C]}$ so any feasible $E\left[\frac{1}{C}\right] \leq \bar{z}$ leads to higher expected profit under the automated contract than under the voluntary contract.

Finally notice that $K \leq 1/C_H$. Therefore $\bar{y} \geq \bar{z}$.

Proof of Theorem 2

$E[1/C]$ may be expressed as $\int_{C_L}^{C_H} \frac{1}{x} g(x) dx$. Consider the approximation of $E[1/C] = \sum_A \frac{1}{x_A} g_A$ where g_A is the discretized version of g . Let us call the support of this discretized distribution $\{x_{-N}, x_{-N+1}, \dots, x_{-1}, x_0, x_1, \dots, x_{N-1}, x_N\}$ where x_0 is the mean around which each pair $\{x_{-A}, x_A\}$ is centered. In particular $x_{-A+1} - x_{-A} = x_A - x_{A-1} = \Delta_A \geq 0$ for all A. To maximize this approximation of $E[1/C]$ we solve

$$\max_g \sum_{A=-N}^N \frac{1}{x_A} g_A$$

subject to

$$\begin{aligned}
g_{-A} &= g_A \forall A \\
\sum_{A=-N}^N &= 1 \\
g_A &\geq 0 \forall A
\end{aligned}$$

where the first constraint imposes symmetry and the second and third constraints ensure that the resulting weights form a proper probability distribution. The knapsack nature of this problem indicates that the program will assign weight either to a single symmetric pair of points or to the mean. The assignment is determined by the ratio of value in the objective to cost in the constraint, $\frac{1/x_A + 1/x_{-A}}{2}$.

The program does not assign weight to the mean if

$$\begin{aligned}
0 &\leq \frac{1/x_1 + 1/x_{-1}}{2} - 1/x_0 \\
&= \frac{\frac{1}{x_0 + \Delta_1} + \frac{1}{x_0 - \Delta_1}}{2} - 1/x_0 \\
&= \frac{x_0}{x_0^2 - \Delta_1^2} - 1/x_0 \\
&= x_0^2 - x_0^2 + \Delta_1^2
\end{aligned}$$

which is satisfied by the assumption that $\Delta_A \geq 0 \forall A$.

The program prefers to weight the pair $\{x_{-n}, x_n\}$ over $\{x_{-n+1}, x_{n-1}\}$ if

$$\frac{1/x_n + 1/x_{-n}}{2} \geq \frac{1/x_{n-1} + 1/x_{-n+1}}{2} = \frac{\frac{1}{x_n + \Delta_n} + \frac{1}{x_{-n} - \Delta_n}}{2}.$$

This condition simplifies to

$$\begin{aligned}\frac{x_n + x_{-n}}{x_n x_{-n}} &\geq \frac{x_n + x_{-n}}{(x_{-n} + \Delta_n)(x_n - \Delta_n)} \\ &= \frac{x_n + x_{-n}}{x_{-n}x_n - x_{-n}\Delta_n + x_n\Delta_n - \Delta_n^2}\end{aligned}$$

which is satisfied as long as $x_n \geq x_{-n} + \Delta_n$ which always holds by assumption that $\Delta_A \geq 0 \forall A$.

We conclude that $E[1/C]$ is increasing in the weight placed on symmetric pairs further from the mean. In other words $E[1/C]$ is increasing in the variance of C . We can generalize this result by taking the limit as $N \rightarrow \infty$ and $\Delta_A \rightarrow 0 \forall A$.

Proof of Theorem 3

We must begin by showing that $W_V^L \geq \Pi_V^L$. From Table 1

$$\begin{aligned}\Pi_V^L &= \frac{a^2 N E \left[\frac{1}{C} \right]^2}{4(2E \left[\frac{1}{C} \right] + bNE \left[\frac{1}{C^2} \right] - \frac{1}{C_H})} \\ &= \frac{a^2 N E \left[\frac{1}{C} \right]^2}{2(2E \left[\frac{1}{C} \right] + bNE \left[\frac{1}{C^2} \right] - \frac{1}{C_H})} - \frac{a^2 E \left[\frac{1}{C} \right]^2 N (bNE \left[\frac{1}{C^2} \right] + 2E \left[\frac{1}{C} \right] - \frac{1}{C_H})}{4(2E \left[\frac{1}{C} \right] + bNE \left[\frac{1}{C^2} \right] - \frac{1}{C_H})^2}. \quad (*)\end{aligned}$$

It follows that $W_V^L \geq \Pi_V^L$ if and only if

$$\begin{aligned} (*) &\leq \frac{a^2 N E \left[\frac{1}{C} \right]^2}{2(2E \left[\frac{1}{C} \right] + bNE \left[\frac{1}{C^2} \right] - \frac{1}{C_H})} - \frac{a^2 N E \left[\frac{1}{C} \right]^2 (bNE \left[\frac{1}{C^2} \right] + E \left[\frac{1}{C} \right])}{4(2E \left[\frac{1}{C} \right] + bNE \left[\frac{1}{C^2} \right] - \frac{1}{C_H})^2} \\ \implies bNE \left[\frac{1}{C^2} \right] + E \left[\frac{1}{C} \right] &\leq bNE \left[\frac{1}{C^2} \right] + 2E \left[\frac{1}{C} \right] - \frac{1}{C_H}\end{aligned}$$

which is always true because $E \left[\frac{1}{C} \right] \geq \frac{1}{C_H}$.

Now that we know that $W_V^L \geq \Pi_V^L$, we need only note that $W_A^L = \Pi_A^L$ to complete the proof. For Part 1, if $\Pi_V^L \geq \Pi_A^L = W_A^L$ then surely $W_V^L \geq W_A^L$. Similarly for Part 2, if $\Pi_A^L = W_A^L \geq W_V^L$ then surely $\Pi_A^L \geq \Pi_V^L$.

Proof of Theorem 4

Plugging w_j^L into $E[q^{LIN}(w_j)]$ in Table 1 yields

$$E[q_V^L(w_V^L)] = \frac{aE\left[\frac{1}{C}\right]^2}{2(2E\left[\frac{1}{C}\right] + bNE\left[\frac{1}{C}^2\right] - \frac{1}{C_H})}$$

$$E[q_A^L(w_A^L)] = \frac{a}{2(E[C] + bN)}$$

Reorganizing demonstrates that $E[q_V^L(w_V^L)] \geq E[q_A^L(w_A^L)]$ if and only if $E[C] - (2 - \frac{1}{C_H E[\frac{1}{C}]}) \frac{1}{E[\frac{1}{C}]} \geq bN \frac{Var(\frac{1}{C})}{E[\frac{1}{C}]^2}$, which is exactly the necessary and sufficient condition for $\Pi_V^L \geq \Pi_A^L$.

Proof of Theorem 5

Let us call $\theta = -C$ where θ is continuous between θ_L and θ_H with probability and cumulative density functions f and F respectively. Given a menu $(w_V(\hat{\theta}), q_V(\hat{\theta}))$ the consumer chooses his payment and quantity by reporting a type $\hat{\theta}$. The consumer with true type θ has value from reporting type $\hat{\theta}$:

$$U(\hat{\theta}, \theta) = w_V^O(\hat{\theta}) + \theta q_V^O(\hat{\theta})^2.$$

The firm must choose the form of the menu $(w_V^O(\hat{\theta}), q_V^O(\hat{\theta}))$ to maximize its profits such that consumers participate (participation constraint: PC) and report their true types (incentive compatibility: IC):

$$\max_{w_V(\cdot), q_V(\cdot)} \int_{\theta_L}^{\theta_H} ((a - bNq_V(\theta))Nq_V(\theta) - w_V(\theta)N)f(\theta)d\theta$$

subject to

$$U(\theta, \theta) \geq U(\hat{\theta}, \theta) \text{ for all } \hat{\theta} \tag{IC}$$

$$U(\theta, \theta) \geq 0 \text{ for all } \theta. \tag{PC}$$

This is a standard adverse selection problem for a continuum of types. We may reduce the continuum of constraints imposed on the firm's problem to the three below.

Lemma 1. *The following three conditions are necessary and sufficient to ensure that the constraints (IC) and (PC) above hold.*

1. $w_V^O(\theta) = U(\theta_L, \theta_L) + \int_{\theta_L}^{\theta} q_V^O(s)^2 ds - \theta q_V^O(\theta)^2$
2. $q_V^O(\theta)$ is non-decreasing in θ
3. $U(\theta_L, \theta_L) \geq 0$.

Proof of Lemma 1

Necessity: Suppose (IC) and (PC) hold. Define $v(\theta) = U(\theta, \theta)$. From (IC) $U_1(\theta, \theta) = 0$ so $\dot{v}(\theta) = U_2(\theta, \theta)$ where the subscript indicates the argument with respect to which a partial derivative is taken. $U_2(\theta, \theta) = q(\theta)^2$ so integrating

$$v(\theta) - v(\theta_L) = \int_{\theta_L}^{\theta} \dot{v}(\theta) = \int_{\theta_L}^{\theta} q(s)^2 ds$$

and then plugging in the expression for $v(\theta)$

$$w(\theta) + \theta q(\theta)^2 - v(\theta_L) - v(\theta_L) = \int_{\theta_L}^{\theta} q(s)^2 ds$$

which, when reorganized, is condition 1. Also by (IC) the taxation principle holds, i.e.

$$(w(), q()) \in \operatorname{argmax} w + \theta q^2.$$

The right hand side of the expression above has increasing differences, i.e. $\frac{\delta^2(w + \theta q^2)}{\delta \theta \delta q} = 2q \geq 0$, so it follows that $q(\theta)$ is non-decreasing in θ .

Sufficiency We want to show that $\theta \in \operatorname{argmax}_{\hat{\theta}} U(\hat{\theta}, \theta)$ and that $q(\theta)$ is non-decreasing in

θ . Differentiate U with respect to $\hat{\theta}$:

$$\frac{\delta U}{\delta \hat{\theta}} = \dot{w}(\hat{\theta}) + 2\theta q(\hat{\theta})\dot{q}(\hat{\theta}).$$

We know $2\theta q(\hat{\theta})\dot{q}(\hat{\theta}) \leq (\geq) 2\hat{\theta} q(\hat{\theta})\dot{q}(\hat{\theta})$ for all $\hat{\theta} > (<) \theta$ if $q(\theta)$ is non-decreasing in θ which implies

$$\frac{\delta U(\hat{\theta}, \theta)}{\delta \hat{\theta}} = \dot{w}(\hat{\theta}) + 2\theta q(\hat{\theta})\dot{q}(\hat{\theta}) \leq (\geq) \dot{w}(\hat{\theta}) + 2\hat{\theta} q(\hat{\theta})\dot{q}(\hat{\theta}) = \frac{\delta U(\hat{\theta}, \hat{\theta})}{\delta \hat{\theta}} = 0$$

for all $\hat{\theta} < (>) \theta$ where the final relationship follows from the assumption that condition (1) holds at $\hat{\theta}$. We have shown that $U(\hat{\theta}, \theta)$ is pseudoconcave in $\hat{\theta}$ with unique maximizer $\hat{\theta} = \theta$, implying (IC) is satisfied.

We may then write the firm's problem as

$$\max_{w_V(\cdot), q_V(\cdot)} \int_{\theta_L}^{\theta_H} ((a - bNq_V(\theta))Nq_V(\theta) - Nw_V(\theta))f(\theta)d\theta \quad (\text{A.1})$$

subject to

$$w_V^O(\theta) = U(\theta_L, \theta_L) + \int_{\theta_L}^{\theta} q_V^O(s)^2 ds - \theta q_V^O(\theta)^2$$

$$\frac{\delta q_V^O(\theta)}{\delta \theta} \geq 0$$

$$U(\theta_L, \theta_L) \geq 0.$$

Plugging the expression for $w_V^O(\theta)$ into (A.1):

$$\max_{q_V(\cdot), U(\theta_L, \theta_L)} \int_{\theta_L}^{\theta_H} \left(aNq_V(\theta) - bN^2q_V(\theta)^2 - N \left(U(\theta_L, \theta_L) + \int_{\theta_L}^{\theta} q_V(s)^2 ds - \theta q_V(\theta)^2 \right) \right) f(\theta) d\theta \quad (\text{A.2})$$

Notice that (A.2) is linearly decreasing in $U(\theta_L, \theta_L)$ so at optimum $U^L(\theta_L, \theta_L) = 0$. Now

let us address the problem of the double integral in (A.2). Integrating by parts:

$$\int_{\theta_L}^{\theta_H} \left(\int_{\theta_L}^{\theta} q_V(s)^2 ds \right) f(\theta) d\theta = \left[\int_{\theta_L}^{\theta} q_V^2(s) ds F(\theta) \right]_{\theta_L}^{\theta_H} - \int_{\theta_L}^{\theta_H} q_V(\theta)^2 F(\theta) d\theta = \int_{\theta_L}^{\theta_H} (1-F(\theta)) q_V(\theta)^2 d\theta$$

which allows us to rewrite the firm's problem as

$$\max_{q_V(\cdot)} \int_{\theta_L}^{\theta_H} \left((aNq_V(\theta) - bN^2q_V(\theta)^2 + \theta Nq_V(\theta)^2) f(\theta) - N(1 - F(\theta))q_V(\theta)^2 \right) d\theta \quad (\text{A.3})$$

subject to

$$\frac{dq_V^Q(\theta)}{d\theta} \geq 0.$$

Maximizing pointwise, we find that the solution to the relaxed problem satisfies its constraint when θ has a non-decreasing failure rate.

A.2. Proofs of Chapter 2 Theorems

Proof of Theorem 6. With a fixed contract the platform chooses p , w , and N . Price can be selected from one of two regions, corresponding to whether demand is served in both demand states or only in the high-demand state: $p < a_l$ and $a_l \leq p < a_h$. We consider each region separately. Suppose $p < a_l$. The platform's expected profit is

$$U = \begin{cases} (p-w)G(w)N & G(w)N \leq a_l - p \\ (p-w)((a_l-p)f_l + G(w)Nf_h) & a_l - p \leq G(w)N \leq a_h - p \\ (p-w)((a_l-p)f_l + (a_h-p)f_h) & a_h - p \leq G(w)N \end{cases}$$

and the utilization of a provider is implicitly defined by $\phi_j = \min\{1, (a_j - p)/NG(\phi_j w)\}$.

The provider's expected profit conditional on joining in period 1 is Π ,

$$\Pi = \begin{cases} \int_0^w G(c)dc & G(w)N \leq a_l - p \\ \int_0^{\phi_l w} G(c)dc f_l + \int_0^w G(c)dc f_h & a_l - p \leq G(w)N \leq a_h - p \\ \int_0^{\phi_l w} G(c)dc f_l + \int_0^{\phi_h w} G(c)dc f_h & a_h - p \leq G(w)N \end{cases}$$

The best contract does not exist exclusively in the first domain of the provider profit function - U strictly increases in N while Π is independent of N , so N must be at least $(a_l - p)/G(w)$.

The optimal solution does not exist exclusively in the third domain of the provider profit function - ϕ_j is decreasing in N , so decreasing N allows w to be decreased, strictly increasing U . So N must be at most $(a_h - p)/G(w)$.

Given the optimal contract is in the second domain of U , the platform's profit is strictly increasing in N . This implies that either the provider profit constraint binds, $c_1 = \Pi$, or the upper bound on the feasible region binds, $NG(w) = a_h - p$. If the former is not true but the latter is, i.e. $c_1 < \Pi$ and $NG(w) = a_h - p$, then the platform's profit is strictly decreasing in w . As $\phi_l w$ is increasing in w , a reduction in w is feasible (because $c_1 < \Pi$), which increases platform profit, which leads to a contradiction. Thus, if the optimal solution has $p < a_l$, then it must be that $a_l - p = N\bar{\phi}_l G(\bar{\phi}_l w)$, which, when substituted into U and the feasible region constraint yields

$$U = (p - w)(a_l - p) \left(f_l + G(w)f_h / (\bar{\phi}_l G(\bar{\phi}_l w)) \right) \quad (\text{A.4})$$

and

$$\bar{p} = (G(w)a_l - \bar{\phi}_l G(\bar{\phi}_l w)a_h) / (G(w) - \bar{\phi}_l G(\bar{\phi}_l w)) \leq p \quad (\text{A.5})$$

As the platform profit (A.4) is concave in p , the optimal price, subject to the constraint (A.5), is $p = \max((a_l + w)/2, \bar{p})$, which satisfies the $p < a_l$ constraint as long as $w' < w < a_l$.

To satisfy the $0 \leq \bar{\phi}_l \leq 1$ constraint, it must be that $w' \leq w \leq w''$. Thus, a search over $w \in [w', \min(w'', a_l)]$ finds the optimal wage.

Suppose $a_l < p < a_h$. The provider joining constraint is $\int_0^w G(c)dc f_h \geq c_1$. The platform's expected profit is

$$U = \begin{cases} (p-w)G(w)N f_h & 0 < G(w)N \leq a_h - p \\ (p-w)(a_h - p)f_h & a_h - p \leq G(w)N \end{cases}$$

If $a_h - p < G(w)N$, then the platform's profit is strictly decreasing in w , so the best fixed contract must satisfy $G(w)N \leq a_h - p$. In this regime, U is strictly increasing in N , so it must be that $G(w)N = a_h - p$. Therefore, $U = (p-w)(a_h - p)f_h$, which is strictly concave in p , so the optimal price is $p = \max\{(a_h + w)/2, a_l\}$. With either price, the platform's profit is strictly decreasing in w , so with the optimal contract the optimal wage is $w = w''$ because that is the wage, by definition, that results in $\Pi = c_1$.

Proof of Theorem 7

To simplify notation and without loss of generality, we assume w_j is paid for participation rather than for service - the w_j wage paid for participation is equivalent to the wage w_j/ϕ_j paid for service. Suppose the platform serves both demand states, i.e. $p < a_l$. The provider joining constraint is

$$\Pi(w) = \sum_j \left(\int_0^{w_j} G(x)dx \right) f_j \geq c_1$$

The platform's profit conditional on a_j is

$$U_j = \begin{cases} (p-w_j)G(w_j)N & NG(w_j) \leq a_j - p \\ p(a_j - p) - w_jG(w_j)N & a_j - p < NG(w_j) \end{cases}$$

In an optimal solution with positive profit, either $N = (a_l - p)/G(w_l)$, or $N = (a_h - p)/G(w_h)$. To explain, for a particular a_j , either $N = (a_j - p)/G(w_j)$ is optimal or $N = 0$ is

optimal. As U_j is linear in N , one of three values for N is selected in the optimal solution: $N = (a_l - p)/G(w_l)$, or $N = (a_h - p)/G(w_h)$, or $N = 0$, but the latter option does not yield a positive profit.

In an optimal solution the provider profit constraint is binding, i.e. $\Pi = c_1$. This can be proven by contradiction. Suppose $\Pi > c_1$ and $(a_l - p)/G(w_l) \leq (a_h - p)/G(w_h)$. As there are only two possible values for N , the profit in the low state is

$$U_l = \begin{cases} (p - w_l)(a_l - p) & N = (a_l - p)/G(w_l) \\ p(a_l - p) - w_l G(w_l)(a_h - p)/G(w_h) & N = (a_h - p)/G(w_h) \end{cases}$$

With either value of N the platform's profit is decreasing in w_l , but because $\Pi > c_1$ allows a reduction in w_l , the current solution cannot be optimal. Suppose $(a_h - p)/G(w_h) < (a_l - p)/G(w_l)$. Now U_h is decreasing in w_h , which contradicts that $\Pi > c_1$ is optimal. Thus, in an optimal solution, $\Pi(w) = c_1$. Define the function $w_h(w_l)$ to return the unique w_h value such that $\Pi(w) = c_1$ for the given w_l . From the implicit function theorem: $\partial w_h / \partial w_l = -q/\gamma(1 - q)$, where for shorthand $q = f_l$, $(1 - q) = f_h$, and $\gamma = G(w_h)/G(w_l)$.

In the optimal solution there is no capacity rationing, i.e. $N = \min \{(a_l - p)/G(w_l), (a_h - p)/G(w_h)\}$.

This can be proven by contradiction. First suppose in the optimal solution $(a_l - p)/G(w_l) < (a_h - p)/G(w_h) = N$. Then U is decreasing in w_l , and so contradicts the initial assumption that the proposed solution is optimal. Second suppose in the optimal solution $(a_h - p)/G(w_h) < (a_l - p)/G(w_l) = N$. Then U is increasing in w_l , which contradicts the assumption that the proposed solution is optimal.

In the optimal solution there is no demand rationing in the low demand state, i.e. $N = (a_l - p)/G(w_l) \leq (a_h - p)/G(w_h)$. To prove this, note that we have already shown that providers are not rationed in the optimal solution, i.e. either $N = (a_l - p)/G(w_l) \leq (a_h - p)/G(w_h)$ or $N = (a_h - p)/G(w_h) \leq (a_l - p)/G(w_l)$. It remains to show that $N = (a_h - p)/G(w_h) < (a_l - p)/G(w_l)$ cannot be optimal. Proof by contradiction. Suppose in the optimal solution

$N = (a_h - p)/G(w_h) < (a_l - p)/G(w_l)$, i.e. $(a_h - p)/(a_l - p) < \gamma$. Differentiate U :

$$\frac{dU}{dw_l} = -q(p - w_l) \frac{(a_h - p)}{\gamma^2} \frac{\partial \gamma}{\partial w_l}$$

Assuming $w_l < p$, it is optimal to increase w_l until the constraint $(a_h - p)/(a_l - p) < \gamma$ binds, which means demand rationing does not occur in the low demand state.

To demonstrate that $w_l < p$ indeed holds in an optimal solution, define \bar{w}_l such that $\frac{a_h}{a_l} = \gamma(\bar{w}_l)$. It follows that $(a_h - p)/(a_l - p) < \gamma$ can only be satisfied for w_l values such that $w_l < \bar{w}_l$, which implies $1 < a_h/a_l < \gamma$, which implies $G(w_l) < G(w_h)$, which implies $w_l < w_h$. For the profit function to be positive, one of these wages must be less than the price, so it follows that in an optimal solution with a positive profit, $w_l < p$.

Given the previous results, the optimal contract is the solution to the following optimization:

$$\begin{aligned} \max_{w_l, p} U &= q(p - w_l)G(w_l)N + (1 - q)(p - w_h)G(w_h)N \\ \text{s.t.} \quad &\gamma = \frac{G(w_h)}{G(w_l)} \leq \frac{a_h - p}{a_l - p} \end{aligned} \tag{A.6}$$

where $w_h = w_h(w_l)$, and $N = (a_l - p)/G(w_l)$. Given that γ is decreasing in w_l and the right hand side of the constraint (A.6) is increasing in p , there exists a p' , which is possibly 0, such that the constraint is satisfied for all $p' \leq p \leq a_l$. As U is concave in p , for a fixed w_l the optimal price is

$$p = \max \left\{ p', \frac{a_l}{2} + \frac{(1 - q)\gamma w_h + q w_l}{2((1 - q)\gamma + q)} \right\} \tag{A.7}$$

In the optimal solution $w_l \leq w'$. This can be proven by contradiction. Assume in the optimal solution $w' < w_l$, which implies $\gamma < 1$, $w_h < w_l$ and $p' = 0$. Hence, the optimal price must be $p = \frac{a_l}{2} + \frac{(1 - q)\gamma w_h + q w_l}{2((1 - q)\gamma + q)}$. Differentiation of the platform's profit function shows

U to be quasiconcave in w_l :

$$\frac{dU}{dw_l} = (a_l - p)(1 - q)(p - w_h) \frac{\partial \gamma}{\partial w_l}$$

which can only be zero if $p = w_h < w_l$, but if the price is less than all wages, the platform's profit cannot be positive, contradicting the assumption that the solution is optimal. Hence a search over w_l in the range $[0, w']$ yields the optimal solution.

Suppose the platform only serves high demand, i.e. $a_l < p < a_h$. This regime is the same as the high price regime in Theorem 1.

Proof of Theorem 8

Suppose $p_l < a_l$. Given wage, w , and recruitment, N , in state a_j platform profit is

$$U_j = \begin{cases} (p_j - w)G(w)N, & G(w)N \leq a_j - p_j \\ (p - w)(a_j - p_j), & a_j - p_j \leq G(w)N \end{cases}$$

In the first region, demand exceeds participation so provider utilization is $\phi_j = 1$. Consequently changes in p have no impact on the providers' profit. Platform profit increases in p , so the platform increases p until $G(w)N = a_j - p_j$. Hence, the optimal policy has $a_j - NG(w) \leq p_j$ for all demand states. The platform's problem is therefore

$$\begin{aligned} \max_{N, p_l, p_h, w} U &= (p_l - w)(a_l - p_l)f_l + (p_h - w)(a_h - p_h)f_h \\ \text{s.t.} \quad &\int_0^{\phi_l w} G(c)dc f_l + \int_0^{\phi_h w} G(c)dc f_h \geq c_1 \\ &a_j - G(w)N \leq p_j \quad \forall j \end{aligned} \tag{A.8}$$

where $\phi_j = \min\{1, (a_j - p_j)/NG(\phi_j w)\}$. Ignoring the constraints, the platform's profit is a concave function of each price with unique maximizers $p_j^* = (a_j + w)/2$, and p_j^* satisfies the constraint (A.8) for all $\bar{N}_j \leq N$, where $\bar{N}_j = (a_j - w)/(2G(w))$. Because $\bar{N}_l \leq \bar{N}_h$, there

are three possible regions for N : $N < \bar{N}_l$, $\bar{N}_l \leq N \leq \bar{N}_h$, $\bar{N}_h < N$.

Show that with the optimal contract, $\bar{N}_l \leq N \leq \bar{N}_h$. Proof by contradiction. Suppose $\bar{N}_h < N$. In this case p_j^* is feasible for $j = l, h$, and there is some provider rationing, i.e. $\phi_j < 1$. The platform's profit is independent of N and decreasing in w for all $w < a_l$ (which ensures that $U > 0$). Therefore, N can be decreased, thereby increasing ϕ_j , which allows the platform to lower w . This increases the platform's profit, contradicting that the policy is optimal. Suppose $N < \bar{N}_l$. Prices are constrained so that demand and participation match, i.e. $p_j = a_j - G(w)N$. For each demand state, $\phi_j = 1$. The platform's profit is concave in N with maximizer $N^* = \left(\sum_{j \in \{l, h\}} a_j f_j - w \right) / 2G(w) \geq \bar{N}_l$, a contradiction.

The optimal price schedule is $p_l = p_l^*$ and $p_h = a_h - G(w)N$, so $\phi_l < 1$ and $\phi_h = 1$. Assuming $w < a_l$, the platform's problem is

$$\begin{aligned} \max_{N, p_l, p_h, w} U &= \frac{(a_l - w)^2}{4} f_l + (a_h - G(w)N - w)G(w)N f_h \\ \text{s.t.} \quad &\int_0^{\phi_l w} G(c)dc f_l + \int_0^w G(c)dc f_h \geq c_1 \\ &\bar{N}_l \leq N \leq \bar{N}_h \end{aligned} \tag{A.9}$$

Ignoring constraint (A.9), the platform's profit is increasing in N over the interval $[\bar{N}_l, \bar{N}_h]$. The platform chooses the largest N such that $\Pi \geq c_1$ and $\bar{N}_h \geq N$. Suppose $\Pi > c_1$ under the optimal solution. Then $N = \bar{N}_h$, prices are $p_h = (a_h - w)/2$ and $p_l = (a_l - w)/2 < a_l$, and the platform's profit is decreasing for all $w < a_l$. Because Π is increasing in w , this contradicts the optimality of $\Pi > c_1$. Thus, under the optimal recruitment, $\Pi = c_1$, which implies $\phi_l = \bar{\phi}_l$ and $N = a_l - w/2\bar{\phi}_l G(\bar{\phi}_l w)$. A search over $w \in [w', \min(a_l, w'')]]$ leads to the optimal w .

If $p_l > a_l$ then the contract is identical to the high demand-only case in the fixed model.

Proof of Theorem 9. Let wages be a fixed commission, β , of price, i.e. $w_j = \beta p_j$. Let \hat{w}_j be the unique wage that matches supply and demand, i.e. $\hat{w}_j = \beta(a_j - NG(\hat{w}_j))$. The

platform's expected profit for $w_j \leq \hat{w}_j$ is $U_j = (1/\beta - 1)w_j NG(w_j)$, which is increasing in w_j . Hence, the optimal wage is at least \hat{w}_j . The platform's expected profit for $w_j > \hat{w}_j$ is $U = (1/\beta - 1)w_j(a_j - w_j/\beta)$, which is concave in w_j . Thus, the profit maximizing wage is for a given a_j is $\max\{\tilde{w}_j, \hat{w}_j\}$, where $\tilde{w}_j = \beta a_j/2$.

Now consider the platform's optimal recruitment for a given commission. The optimal wage schedule is a function of recruitment: $\hat{w}_j \leq \tilde{w}_j$ if and only if $a_j/2 \leq NG(\hat{w}_j)$ where differentiation shows that $NG(\hat{w}_j)$ is an increasing function of N . Define $\bar{N}_j > 0$ to be the unique recruitment threshold for which $\hat{w}_j < \tilde{w}_j$ if and only if $\bar{N}_j < N$ and define provider utilization given wage \tilde{w}_j to be $\tilde{\phi}_j = a_j/2NG(\tilde{\phi}_j\tilde{w}_j)$. Then expected profit of a provider for a given a_j is

$$\Pi_j = \begin{cases} \int_0^{\tilde{w}_j\tilde{\phi}_j} G(c)dc, & \bar{N}_j < N \\ \int_0^{\hat{w}_j} G(c)dc, & \bar{N}_j \geq N \end{cases}$$

Notice $\tilde{w}_j\tilde{\phi}_j$ is a decreasing function of N , so Π_j is a monotonically decreasing function of N . In contrast, the platform's expected profit from a realization a_j is a weakly increasing function of N :

$$U_j = \begin{cases} (1 - \beta)a_j^2/4, & \bar{N}_j < N \\ (1/\beta - 1)\hat{w}_j(a_j - \hat{w}_j/\beta), & \bar{N}_j \geq N \end{cases}$$

It follows that the platform chooses recruitment so that $\Pi = \sum_{j \in \{l, h\}} \Pi_j f_j = c_1$.

It remains to search over β . Because w_j is decreasing in both N and β , we may find an lower bound on β from $\max_j w_j(N = 0) = \hat{w}_h(N = 0) = \beta a_h \geq w'$. Search for the profit maximizing commission on the interval $[w'/a_h, 1]$.

Proof of Theorem 10. Suppose the platform selects a price p_j and a wage w_j for each demand state a_j to maximize the system's profit (the sum of platform and provider surplus). Although the platform makes two decisions for each demand state, it is possible to reduce this to a single decision because it is never optimal to choose a price/wage combination such that demand doesn't exactly match supply: if demand exceeds supply, system profits

can be increased by raising the price; and if demand is less than supply, system profit can be increased by decreasing the wage. Hence, for any demand state a_j , the price and wage selected must satisfy $NG(w_j) = a_j - p_j$. Let $S_j(p_j(w_j), w_j)$ be the system's expected profit given a wage and demand realization:

$$S_j(p_j(w_j), w_j) = (a_j - p_j)(p_j - w_j) + N \int_0^{w_j} G(c) dc = NG(w_j)(a_j - NG(w_j) - w_j) + N \int_0^{w_j} G(c) dc$$

The system's expected profit, including the cost of having N providers join, is $S(w_j, N) = S_l(w_l)f_l + S_h(w_h)f_h - Nc_1$. Because $S_j(w_j)$ is quasi-concave, there exists a unique w_j^* that maximizes system profit for each demand state a_j : $w_j^* + 2NG(w_j^*) = a_j$, which is decreasing in N . Changing N affects system surplus:

$$\begin{aligned} \frac{dS}{dN} &= \sum_j \left(G(w_j^*)(a_j - 2NG(w_j^*) - w_j^*) + \int_0^{w_j^*} G(c) dc \right) f_j - c_1 = \sum_j \left(\int_0^{w_j^*} G(c) dc \right) f_j - c_1 \\ \frac{\partial^2 S}{\partial N^2} &= \sum_j \frac{\partial w_j^*}{\partial N} G(w_j^*) f_j < 0 \end{aligned}$$

Thus, system profit is concave in N and there exists a unique N^* that maximizes system profit. With $N = N^*$ providers earn their minimum profit, i.e. $\Pi(N^*) = c_1$. It follows that the system optimal solution is also the contract that maximizes the platform's profit subject to $\Pi \geq c_1$. Finally, a series of substitutions yields $p_j^* = w_j^*(1 + NG(w_j^*)/w_j^*)$.

The optimal contract yields a bound on the largest feasible provider reservation price. As c_1 becomes large, the platform extracts all surplus from consumers by charging $p_j^* \rightarrow a_j$ and passes all profit to providers via $w_j^* \rightarrow a_j$. The platform earns weakly positive profit and providers earn $\sum_{j \in \{l, h\}} \int_0^{a_j} G(c) dc f_j$. Then the largest c_1 for which the platform can feasibly operate is

$$\bar{c}_1 = \sum_{j \in \{l, h\}} \int_0^{a_j} G(c) dc f_j. \quad (\text{A.10})$$

Proof of Theorem 11

1. As c_1 approaches \bar{c}_1 , the optimal contract parameters behave as follows: $N^o \rightarrow 0$,

$w_j^o \rightarrow a_j$, $p_j^o \rightarrow a_j$. Similarly, under the commission contract, $N \rightarrow 0$. Consequently, $N < \bar{N}_j \forall j$, so $w_j = \beta(a_j - NG(w_j)) \rightarrow \beta a_j$ and providers are fully utilized. To maintain feasibility, the platform must choose $\beta \rightarrow 1$, which recovers the optimal contract.

2. As c_1 approaches 0, the optimal contract parameters behave as follows: $N^o \rightarrow \infty$, $w_j^o \rightarrow 0$, $p_j^o \rightarrow a_j/2$. Similarly, under the commission contract, the payment to providers approaches zero and $N \rightarrow \infty$. Consequently, $N > \bar{N}_j \forall j$, so $w_j = a_j\beta/2$, which implies $p_j = a_j/2$. $N \rightarrow \infty$ causes $\bar{\phi}_j \rightarrow 0$, so the per participation payment $(a_j\bar{\phi}_j\beta)/2 \rightarrow 0$, which recovers the optimal contract.

A.2.1. Fixed c_2 Model

This part derives the results for the “fixed c_2 ” model. The fixed c_2 model is identical to the main model with the one exception that the participation cost for all providers in period 2 is fixed at c_2 . Assume $c_2 < a_l$ (so that the low demand state is viable), and $c_1 < f_h(a_h - a_l) + a_l - c_2$ (so the optimal policy can earn positive profit). For notational convenience, let $\bar{a} = f_l a_l + f_h a_h$ and $\hat{c} = c_2 + c_1/f_h$.

The key difference between the fixed c_2 model and the main model is how the providers’ participation cost varies in the number of joining providers, N , who can potentially participate in period 2. In the fixed c_2 model each participating provider incurs cost c_2 independent of N . In the main model, because provider participation costs are heterogeneous and independent across providers, holding the number of participating providers constant, their average participation cost decreases in N . For example, the average participation cost of the 100 providers with the lowest realized participation costs is lower if they are selected from a set of $N = 1,000$ providers rather than a set of $N = 200$ providers. Consequently, in the main model, if the joining cost is small, it is in the interest of the platform to have a large number of providers join (i.e., large N) because that ensures their average participation cost will be low among the providers who actually participate. This does not occur in the

fixed c_2 model, which means the platform has a lower incentive to increase the number of joining providers in that model. Thus, the two models behave differently for low values of c_1 .

Participation equilibrium

When there is heterogeneity in participation costs (as in the main model) the equilibrium participation strategies can be described as a threshold - a provider participates if his/her participation cost is less than the threshold, otherwise they don't participate. In contrast, when participation costs are fixed and common across all providers, providers use a mixed strategy.

If demand exceeds the number of providers, then all providers participate as long as the wage they receive exceeds c_2 (the participation constraint).

If demand is less than the number of providers in a given state, i.e., $a_j - p_j < N$, then providers must decide whether to participate or not. Let ϕ be the probability a provider participates. In equilibrium

$$\phi = \min \left\{ \left(\frac{w_j}{c_2} \right) \left(\frac{a_j - p_j}{N} \right), 1 \right\}$$

To explain, if

$$c_2 \leq w_j \left(\frac{a_j - p_j}{N} \right)$$

then all providers participate, i.e., $\phi = 1$, even though some are rationed because the participation constraint is not binding. Otherwise, only a fraction of providers participate and they all earn 0 profit (i.e., the participation constraint is binding). Provider earnings are thus

$$\pi_j = \max \left\{ w_j \frac{a_j - p_l}{N} - c_2, 0 \right\}$$

As in the main model, it is costly for the platform and the system to have excessive participation - the platform can increase its wage, but this might have zero impact on the providers' earnings because they squander the higher wages with excessive participation that leads to providers with less than full utilization. Thus, as in the main model, the platform has the challenge of providing sufficient capacity when demand is high but not too much capacity when demand is low.

Fixed price

Suppose the platform charges the same price, p , in both demand states and compensates providers such that their joining constraint is binding, as occurs in the fixed contract in the main model. Hence, the platform's costs include c_2 for each participating provider and c_1 for each joining provider. This compensation can be achieved with a fixed wage for service (equal to c_2 , so that all providers who participate receive demand) and a fixed salary for joining the platform (equal to c_1 , to ensure the joining constraint is satisfied). Alternatively, it can be achieved with a base wage that is paid in all states (equal to c_2) and a supplemental wage that is paid in the high demand state (equal to c_1/f_h) - in the main model that would be analogous to the dynamic wage contract. It cannot be achieved with a fixed wage: unlike in the main model, with a fixed wage some surplus is lost due to idle providers who nevertheless incur participation costs.

Only three possible strategies can be optimal for the platform, which we describe as “poor service”, “poor utilization”, and “only high demand”. While optimal prices, profits and surplus are derived in closed form for each case, simple conditions do not exist to determine which of the three strategies is optimal for a particular setting. Nevertheless, it is trivial to compare them numerically.

Case 1: Poor service: $p < a_l$, $N = a_l - p$

With this strategy demand and price are the same in both states. Hence, there is demand

rationing in the high demand state. The platform's profit is

$$U_f(p) = (a_l - p)(p - c_1 - c_2)$$

The optimal price is

$$p_f^* = \frac{a_l + c_1 + c_2}{2}$$

and the platform's resulting profit is

$$U_f(p_{f1}^*) = \frac{(a_l - c_1 - c_2)^2}{4}$$

Consumer surplus is

$$S_f(p) = \frac{1}{2}f_l(a_l - p)^2 + \frac{1}{2}f_h(a_h - p)(a_l - p)$$

which yields

$$S_f^* = S_f(p_{f1}^*) = \frac{1}{8} \left((\bar{a} - c_1 - c_2)^2 - f_h^2 (a_h - a_l)^2 \right)$$

Case 2: Poor utilization: $p < a_l$, $N = a_h - p$

With this strategy the platform matches capacity to demand in the high state. Hence, there is capacity rationing in the low demand state. The platform's profit is

$$U_f(p) = f_l(a_l - p)(p - c_2) + f_h(a_h - p)(p - c_2) - (a_h - p)c_1$$

The optimal price is

$$p_f^* = \min \left\{ \frac{\bar{a} + c_1 + c_2}{2}, a_l \right\}$$

and the platform's resulting profit is (assuming $p_{f2}^* < a_l$)

$$U_f(p_f^*) = \frac{1}{4} (\bar{a} - c_1 - c_2)^2 - f_l(a_h - a_l)c_1$$

Consumer surplus is

$$S_f(p) = \frac{1}{2} \left(f_l (a_l - p)^2 + f_h (a_h - p)^2 \right)$$

which yields

$$S_f^* = S_f(p_f^*) = \frac{1}{8} (\bar{a} - c_1 - c_2)^2 + \frac{1}{2} (f_l a_l^2 + f_h a_h^2 - \bar{a}^2)$$

Case 3: Only high demand: $a_l < p$, $N = a_h - p$

With this strategy the platform sets price sufficiently high to only serve demand in the high state. The platform's profit is

$$U_f(p) = f_h (a_h - p)(p - c_2) - (a_h - p)c_1$$

The optimal price is

$$p_f^* = \max \left\{ \frac{a_h + \hat{c}}{2}, a_l \right\}$$

and the platform's resulting profit is (assuming $a_l < p_f^*$)

$$U_f(p_f^*) = \frac{f_h}{4} (a_h - \hat{c})^2$$

Consumer surplus is

$$S_f(p) = \frac{1}{2} f_h (a_h - p) (a_h - p)$$

which yields

$$S_f^* = S_f(p_f^*) = \frac{1}{8} f_h (a_h - \hat{c})^2$$

Optimal Contract

The optimal contract specifies a price, p_j , and a wage, w_j , for each state. For ease of comparison with the fixed commission contract, the optimal contract can also be defined in terms of a state-specific price and commission, $\beta_j = w_j/p_j$, as is done in this section.

We begin with several properties that hold for the optimal contract:

- The number of joining providers, N , must equal the highest realized demand. If N exceeds the highest demand, then it can be reduced and profits weakly increase for both providers and the platform. If N is less than the highest realized demand, then in that state an increase in price increases both the platform's and the provider's profit.
- Both demand states are served - the platform and providers are always able to earn some profit from a demand state conditional on the actions in the other demand state.
- Demand in the high state cannot be strictly lower than demand in the low state, i.e., $a_h - p_h < a_l - p_l = N$ is not optimal. If that were to occur, then the optimal prices are $p_l = (a_l + \hat{c})/2$ and $p_h = (a_h + c_2)/2$. But with those prices, demand in the high demand state is greater than demand in the low demand state, contradicting the original assumption.

Two possibilities for the optimal contract remain (all serving demand in both states). If $c_1 < f_h(a_h - a_l)$ then the optimal contract yields $a_l - p_l < a_h - p_h = N$, otherwise with the optimal contract $a_l - p_l = a_h - p_h = N$. In the later case, the optimal contract can be implemented with a single commission rate, i.e., $\beta_l = \beta_h = \beta$.

Case 1: $a_l - p_l < a_h - p_h$

Assume with the optimal contract $a_l - p_l < a_h - p_h$ (meaning that there can be capacity rationing in the low demand state), all participating providers in the low demand state are fully utilized (thereby there is no loss of system value due to idle providers), providers earn zero profit in the low-demand state, all providers participate in the high demand state and providers earn c_1/f_h in the high demand state. If such a contract is found, then it is an optimal contract for the platform because it maximizes system surplus while the providers' joining constraint is binding (i.e., providers earn their minimum requirement).

To construct the desired contract, the participation constraint in the low-demand state should bind, i.e., $\beta_l p_l = c_2$ so that precisely $a_l - p_l$ providers participate (and all of them are fully utilized). The joining constraint is then

$$c_1 \leq \Pi = f_l(\beta_l p_l - c_2) + f_h(\beta_h p_h - c_2)$$

or

$$\hat{c} \leq \beta_h p_h$$

which should also bind, i.e., $\beta_h = \hat{c}/p_h$.

Under these conditions, the platform's profit can be written as

$$U_o = f_l(p_l - c_2)(a_l - p_l) + f_h(p_h - \hat{c})(a_h - p_h)$$

The optimal prices and commission rates are

$$p_l^o = \frac{a_l + c_2}{2}$$

$$p_h^o = \frac{a_h + \hat{c}}{2}$$

$$\beta_l^o = \frac{2c_2}{a_l + c_2}$$

$$\beta_h^o = \frac{2\hat{c}}{a_h + \hat{c}}$$

The resulting profit is

$$U_o = \frac{f_l}{4}(a_l - c_2)^2 + \frac{f_h}{4}(a_h - \hat{c})^2$$

The original assumption that demand in the high state is greater than demand in the low state requires

$$a_l - p_l < a_h - p_h$$

or

$$c_1 < f_h(a_h - a_l)$$

Note that β_h^o is increasing in c_1 , while β_l^o is constant, and $\beta_h^o < \beta_l^o$ for low values of c_1 . In particular, $\beta_l^o \leq \beta_h^o$ when

$$\left(\frac{c_2}{a_l}\right) f_h(a_h - a_l) \leq c_1 < f_h(a_h - a_l)$$

Case 2: $a_l - p_l = a_h - p_h$

In the second form of the optimal contract $a_l - p_l = a_h - p_h$, meaning that there isn't capacity rationing in either state. The platform operates under several constraints

$$\begin{aligned} c_2 &\leq \beta_l p_l \\ c_2 &\leq \beta_h p_h \\ c_1 &\leq f_l(\beta_l p_l - c_2) + f_h(\beta_h p_h - c_2) \\ p_l &= p_h - (a_h - a_l) \\ p_l &\leq a_l \\ p_h &\leq a_h \end{aligned}$$

The first two are participation constraints, the third is the joining constraint, the fourth ensures that demand is equal in both states and the last two ensure demand is served in both states. By removing redundant constraints, the above set can be written as

$$\begin{aligned} c_2/\beta_l + (a_h - a_l) &\leq p_h \\ c_2/\beta_h &\leq p_h \\ \frac{c_1 + c_2 + f_l \beta_l (a_h - a_l)}{f_l \beta_l + f_h \beta_h} &\leq p_h \\ p_h &\leq a_h \end{aligned}$$

The first two are the participation constraints, the third is the joining constraint and the fourth ensures demand is positive in the high state.

Assume the platform can find commission rates such that the participation and joining constraints are satisfied and the joining constraint is binding. In that case, the platform's

profit is

$$\begin{aligned} U_o &= f_l (p_l - c_2) (a_l - p_l) + f_h (p_h - c_2) (a_h - p_h) - c_1 (a_h - p_h) \\ &= (a_h - p_h) (p_h - c_1 - c_2 - (1 - f_h) (a_h - a_l)) \end{aligned}$$

The resulting optimal prices are

$$\begin{aligned} p_h^o &= \frac{1}{2} (a_h + c_1 + c_2 + (1 - f_h) (a_h - a_l)) \\ p_l^o &= \frac{1}{2} (a_l + c_1 + c_2 - f_h (a_h - a_l)) \end{aligned}$$

and the resulting profit for the platform is

$$U_o = \frac{1}{4} (f_l a_l + f_h a_h - c_1 - c_2)^2$$

It is possible to achieve the optimal profit with different commission rates in both states because in both states the commission rate serves to transfer profit to the providers (so that the joining constraint is satisfied) without the concern of idle providers (because N equals demand in both states). However, it is also possible to achieve the optimal profit with a single commission rate. Assuming $\beta_l = \beta_h = \beta$, the joining constraint and the solution for p_h^o yields the optimal commission:

$$\beta^o = \frac{2(c_1 + c_2)}{\bar{a} + (c_1 + c_2)}$$

Finally, the $p_h < a_h$ constraint implies $c_1 + c_2 < \bar{a}$, which can be written as

$$c_1 < f_h (a_h - a_l) + a_l - c_2$$

Consumer Surplus

Consumer surplus is

$$S_o(p_l, p_h) = \frac{1}{2}f_l(a_l - p_l)^2 + \frac{1}{2}f_h(a_h - p_h)^2$$

which yields with the optimal prices

$$S_o = S_o(p_l^o, p_h^o) = \begin{cases} s_l(c_1) & c_1 < f_h(a_h - a_l) \\ s_h(c_1) & f_h(a_h - a_l) < c_1 \end{cases}$$

where

$$s_l(c_1) = \frac{1}{8} \left(f_l(a_l - c_2)^2 + f_h(a_h - \hat{c})^2 \right)$$

and

$$s_h(c_1) = \frac{1}{8} (\bar{a} - c_1 - c_2)^2$$

Proof of Proposition 1

Each subsection compares consumer surplus with one of the three fixed-price strategies to the optimal contract. Switching from the fixed price contract to the optimal contract decreases consumer surplus only when the fixed-price strategy implements the poor utilization version (i.e., capacity is rationed). Otherwise, switching to the optimal contract increases consumer surplus.

Case 1: Poor service

Consumer surplus with the optimal pricing strategy is greater than a single-price strategy with demand rationing.

The optimal surplus is greater when $f_h(a_h - a_l) < c_1$ because clearly

$$S_f^* < s_h(c_1).$$

Now assume $c_1 \leq f_h(a_h - a_l)$. Define

$$y(c_1) = \frac{1}{8} (\bar{a} - c_1 - c_2)^2$$

Clearly

$$S_f^* < y(c_1)$$

so it is sufficient to show that

$$y(c_1) \leq s_l(c_1)$$

To see this, note that

$$y(f_h(a_h - a_l)) = s_l(f_h(a_h - a_l))$$

and

$$0 > \frac{\partial y(c_1)}{\partial c_1} > \frac{\partial s_l(c_1)}{\partial c_1}$$

Case 2: Poor utilization

Consumer surplus with the optimal contract is less than a fixed-price contract with capacity rationing.

The result, $S_o < S_f^*$, clearly applies when $f_h(a_h - a_l) < c_1$.

Now assume $c_1 \leq f_h(a_h - a_l)$. First note that

$$s_l(0) < S_f^*(0)$$

and

$$s_l(f_h(a_h - a_l)) < S_f^*(f_h(a_h - a_l))$$

To confirm that $s_l(c_1) < S_f^*(c_1)$, note that

$$0 > \frac{\partial S_f^*}{\partial c_1} > \frac{\partial s_l(c_1)}{\partial c_1}$$

Case 3: High demand only

Straightforward algebra and the constraint $c_1 + c_2 < a_l$ demonstrates that the optimal contract increases consumer surplus relative to the fixed contract that only serves high demand.

Fixed commission

Let β be the fixed commission the platform pays in both states.

It is possible to quickly rule out several cases for the fixed commission contract. It is not optimal to serve only low demand: Given prices and β for a low demand case, it is always possible to find a p_h that generates some high-state demand and therefore higher profit. It is also intuitive that it is not optimal to choose parameters such that low demand exceeds high demand. Three possible versions of the optimal fixed commission contract remain.

Case 1: Only high demand

Suppose the platform chooses a commission rate such that it abandons the low demand market. For example, if $\beta < c_2/a_l$ then $a_l < p_l$ is required for participation, which is clearly not profitable. This may occur because the joining cost is sufficiently low, which means the platform's optimal commission is small.

The platform profit is then

$$\begin{aligned}
 U_\beta &= f_h (1 - \beta) p_h (a_h - p_h) \\
 \text{s.t.} \quad &c_2 \leq \beta p_h \\
 &p_h \leq a_h \\
 &c_1 \leq \Pi
 \end{aligned}$$

The combination of the constraints yields

$$\hat{c}/\beta \leq p_h \leq a_h$$

The optimal price and commission rates are

$$\begin{aligned}
 p_h^* &= \frac{a_h + \hat{c}}{2} \\
 \beta^* &= \frac{2\hat{c}}{a_h + \hat{c}}
 \end{aligned}$$

and the resulting platform profit is

$$U_\beta = f_h \frac{(a_h - \hat{c})^2}{4}$$

For the platform to not be willing to participate in the low demand market, it must be that the price that induces participation is so high that it also yields zero demand, i.e.,

$$a_l \leq \frac{c_2}{\beta^*}$$

which can be written as

$$c_1 \leq f_h c_2 \left(\frac{a_h}{2a_l - c_2} - 1 \right)$$

Hence, the fixed commission contract ignores low demand only when the joining cost is sufficiently low.

Case 2: High demand exceeds low demand

There are two versions of the contract that have high demand exceeding low demand. In the first, the joining constraint is binding. If this contract exists, it is the optimal fixed commission contract. In the second, the joining constraint does not bind. The first is analytically tractable but the second is not. Nevertheless, the first provides a lower bound on the platform's fixed commission profit.

Subcase 1: Binding joining constraint

Let's derive a fixed commission contract in which there is no welfare loss in the low demand state and the joining constraint is binding, i.e.,

$$\beta p_l = c_2$$

and

$$\hat{c} = \beta p_h$$

If this contract exists, conditional that demand is served in both states, then it is the optimal fixed commission contract because surplus is maximized and the provider's profit is minimized.

The $p_l \leq a_l$ constraint implies

$$\frac{c_2}{a_l} \leq \beta$$

The $p_h \leq a_h$ constraint implies

$$\frac{\hat{c}}{a_h} \leq \beta$$

The $a_l - p_l < a_h - p_h$ constraint implies

$$\frac{c_1}{f_h(a_h - a_l)} < \beta$$

It can be shown that the first constraint, $p_l \leq a_l$, is the most restrictive of the three if

$$c_1 < \left(\frac{c_2}{a_l} \right) f_h (a_h - a_l)$$

and the third constraint, $a_l - p_l < a_h - p_h$, is the most restrictive of the three if

$$\left(\frac{c_2}{a_l} \right) f_h (a_h - a_l) < c_1$$

The platform's profit can be written as

$$\begin{aligned} U_\beta &= f_l (c_2/\beta - c_2) (a_l - c_2/\beta) + f_h (\hat{c}/\beta - \hat{c}) (a_h - \hat{c}/\beta) \\ &= \left(\frac{1-\beta}{\beta} \right) \left(f_l c_2 a_l + f_h \hat{c} a_h - \frac{f_l c_2^2 + f_h \hat{c}^2}{\beta} \right) \end{aligned}$$

The platform's profit function is quasi-concave in β and the optimal β is

$$\beta^* = \frac{2 (f_l c_2^2 + f_h \hat{c}^2)}{f_l c_2 a_l + f_h \hat{c} a_h + f_l c_2^2 + f_h \hat{c}^2}$$

The resulting profit is

$$\begin{aligned} U_\beta &= \frac{(f_l c_2 a_l + f_h \hat{c} a_h - f_l c_2^2 - f_h \hat{c}^2)^2}{4(f_l c_2^2 + f_h \hat{c}^2)} \\ &= \frac{(f_l c_2 (a_l - c_2) + f_h \hat{c} (a_h - \hat{c}))^2}{4(f_l c_2^2 + f_h \hat{c}^2)} \end{aligned}$$

Note that as $f_h \rightarrow 1$ then β^* approaches the optimal contract high-demand commission, β_h^o . And as $f_h \rightarrow 0$ then β^* approaches the optimal contract low-demand commission, β_l^o . Thus, $\min \{\beta_l^o, \beta_h^o\} \leq \beta^* \leq \max \{\beta_l^o, \beta_h^o\}$. As with the optimal contract, there are no idle providers, providers earn zero in the low-demand state, and providers earn c_1/f_h in the high-demand state. However, this fixed commission contract is not optimal because either p_l is too low and p_h is too high (when $\beta_l^o < \beta_h^o$) or p_l is too high and p_h is too low (when $\beta_h^o < \beta_l^o$).

If

$$c_1 = \left(\frac{c_2}{a_l} \right) f_h(a_h - a_l)$$

then

$$\beta^* = \frac{2c_2}{a_l + c_2}$$

which yields the optimal profit.

If

$$c_1 < \left(\frac{c_2}{a_l} \right) f_h(a_h - a_l)$$

then the constraint to satisfy is $p_l \leq a_l$, which is written as

$$\frac{c_2}{a_l} \leq \beta^*$$

If

$$\left(\frac{c_2}{a_l} \right) f_h(a_h - a_l) < c_1$$

then it must be that

$$\frac{c_1}{f_h(a_h - a_l)} < \beta^* < \frac{2\hat{c}}{a_h}$$

The lower bound is needed to satisfy $a_l - p_l < a_h - p_h$. The upper bound is needed to ensure that it is optimal for the platform to make the joining constraint bind with β^* . To explain, for a fixed β the optimal p_h is $a_h/2$, whereas the p_h that makes the joining condition bind is \hat{c}/β^* . Hence, for the joining constraint to bind in the optimal solution it must be that $a_h/2 \leq \hat{c}/\beta^*$.

Subcase 2: Non-binding joining constraint

Let's derive a fixed commission contract in which there is no welfare loss in the low demand state and the joining constraint is NOT binding, i.e.,

$$\beta p_l = c_2$$

and

$$\hat{c} < \beta p_h$$

The platform may choose this contract because it does not want to abandon the low demand state, but doing so requires a relatively high commission to ensure participation, and that high commission leaves providers with more than enough surplus in the high demand state to cover their joining constraint.

The platform's profit is

$$U_\beta = f_l (c_2/\beta - c_2) (a_l - c_2/\beta) + f_h (1 - \beta) p_h (a_h - p_h)$$

The optimal p_h is $a_h/2$, leaving the platform profit to be

$$U_\beta = (1 - \beta) (f_l (c_2/\beta) (a_l - c_2/\beta) + f_h (a_h^2/4))$$

It is possible to show that U_β is quasi-concave in β , but a closed form solution is messy:

$$\frac{\partial U_\beta}{\partial \beta} = \frac{(2 - \beta) c_2^2 f_l - f_l a_l \beta c_2 - f_h (a_h^2/4) \beta^3}{\beta^3}$$

Case 3: High demand equals low demand

This contract mimics the optimal contract with equal demand in both states.

The platform's profit is

$$\begin{aligned} U_\beta &= f_l (p_l - c_2) (a_l - p_l) + f_h (p_h - c_2) (a_h - p_h) - c_1 (a_h - p_h) \\ &= (a_h - p_h) (p_h - c_1 - c_2 - (1 - f_h) (a_h - a_l)) \end{aligned}$$

The resulting optimal prices are

$$\begin{aligned} p_h^* &= \frac{1}{2} (a_h + c_1 + c_2 + (1 - f_h) (a_h - a_l)) \\ p_l^* &= \frac{1}{2} (a_l + c_1 + c_2 - f_h (a_h - a_l)) \end{aligned}$$

The condition $p_h \leq a_h$ and the joining constraint require

$$\beta f_h (a_h - a_l) \leq c_1 \leq f_h (a_h - a_l) + a_l - c_2$$

The optimal profit is

$$U_\beta = \frac{(f_l(a_l - c_2) + f_h(a_h - \hat{c}))^2}{4}$$

The optimal fixed commission is

$$\beta^* = \frac{2(c_1 + c_2)}{\bar{a} + (c_1 + c_2)}$$

As β^* is concave and increasing in c_1 , there exists a \hat{c}_1 such that the condition for this contract, $\beta f_h (a_h - a_l) \leq c_1$, is satisfied for all $\hat{c}_1 \leq c_1$. Note that $\hat{c}_1 < f_h (a_h - a_l)$.

Proof of Proposition 2

Our interest is to derive a lower bound on the performance of the fixed commission contract relative to the optimal contract. To begin, there are conditions in which the fixed commission is optimal. For example, if

$$c_1 = \left(\frac{c_2}{a_l} \right) f_h (a_h - a_l)$$

then

$$\beta^* = \beta_l^o = \beta_h^o.$$

In this case the parameter values align such that in the optimal contract it turns out that the fixed commission rates are identical in the two states. We suspect a comparable result

would not extend to the case of multiple demand states. So this is viewed as a special case with less interest (i.e., less generality).

The fixed commission contract is also optimal if

$$f_h(a_h - a_l) < c_1 \leq f_h(a_h - a_l) + a_l - c_2$$

To explain, these are situations in which the joining cost is sufficiently high that even with the optimal contract it is not desirable to serve more customers in the high demand state (though they are charged a higher price). Consequently a single commission rate is sufficient to maximize the platform's profit. This is likely to be generally true even with multiple demand states.

Our interest is to derive a lower bound on the ratio of fixed commission to optimal platform profits, U_β/U_o . Thus, we assume $c_1 < f_h(a_h - a_l)$ in the subsequent discussion because those are the cases in which $U_\beta < U_o$.

In the fixed c_2 model it is possible to evaluate the profit functions for the entire parameter space. We observe that the ratio U_β/U_o is minimized with either $c_2 \rightarrow 0$ or $c_1 \rightarrow 0$. Therefore, we derive the bound for those two cases. Furthermore, because we do not have a closed form solution for U_β when the high demand exceeds low demand but the joining constraint is not binding, we use the comparable solution with a binding joining constraint as a lower bound. Hence, for those situations a derived lower bound is not tight. We observe in most situations our bound is indeed tight.

We derive the lower bound in two cases, $c_2 \rightarrow 0$ or $c_1 \rightarrow 0$. Consider the $c_2 = 0$ case first.

With the optimal policy,

$$U_o(c_2 = 0) = \frac{f_l}{4} a_l^2 + \frac{f_h}{4} (a_h - c_1/f_h)^2$$

With the fixed commission contract and only-high demand or with $0 < a_l - p_l < a_h - p_h$,

$$U_\beta (c_2 = 0) = f_h \frac{(a_h - c_1/f_h)^2}{4}$$

With the fixed commission contract and $0 < a_l - p_l = a_h - p_h$:

$$U_\beta (c_2 = 0) = \frac{(f_l a_l + f_h (a_h - c_1/f_h))^2}{4}$$

The fixed commission profit is

$$U_\beta = \max \left\{ \frac{f_h (a_h - c_1/f_h)^2}{4}, \frac{(f_l a_l + f_h (a_h - c_1/f_h))^2}{4} \right\} \quad (\text{A.11})$$

For notational convenience, let $x = a_h - c_1/f_h$. Note that the constraint $c_1/f_h < a_h - a_l$ implies $a_l < x$. The first term of A.11 is the optimal fixed commission, when

$$a_l \left(\frac{1 - f_h}{\sqrt{f_h} - f_h} \right) < x$$

The profit ratio is

$$\frac{U_\beta}{U_o} = \frac{\frac{f_h (a_h - c_1/f_h)^2}{4}}{\frac{f_l a_l^2}{4} + \frac{f_h}{4} (a_h - c_1/f_h)^2}$$

which can be written as

$$\frac{U_\beta}{U_o} = \frac{f_h x^2}{f_l a_l^2 + f_h x^2}$$

which is increasing in x . Hence the profit ratio is minimized at the lower bound of the feasible region for x . Simplification of the lower bounds yields

$$\min \left\{ \frac{U_\beta}{U_o} \right\} = \frac{1}{2} \left(1 + \sqrt{f_h} \right). \quad (\text{A.12})$$

The second term in A.11 is the optimal fixed commission profit when

$$x < a_l \left(\frac{1 - f_h}{\sqrt{f_h} - f_h} \right)$$

and the ratio of the fixed commission profit to the optimal profit is

$$\frac{U_\beta}{U_o} = \frac{\frac{(f_l a_l + f_h (a_h - c_1 / f_h))^2}{4}}{\frac{f_l}{4} a_l^2 + \frac{f_h}{4} (a_h - c_1 / f_h)^2}$$

which can be written as

$$\frac{U_\beta}{U_o} = \frac{(f_l a_l + f_h x)^2}{f_l a_l^2 + f_h x^2}$$

which is decreasing in x . Hence the minimum ratio occurs at the upper bound of x . After simplification, the lower bound matches (A.12). Thus, (A.12) is a lower bound for the case $c_2 = 0$.

Now let's consider the $c_1 = 0$ case. With the optimal policy,

$$U_o(c_1 = 0) = \frac{f_l}{4} (a_l - c_2)^2 + \frac{f_h}{4} (a_h - c_2)^2$$

With the fixed commission contract and only-high demand

$$U_\beta(c_1 = 0) = f_h \frac{(a_h - c_2)^2}{4}$$

With the fixed commission contract and $0 < a_l - p_l < a_h - p_h$ or $0 < a_l - p_l = a_h - p_h$,

$$U_\beta(c_1 = 0) = \frac{(f_l(a_l - c_2) + f_h(a_h - c_2))^2}{4}$$

Thus, the fixed commission profit is

$$U_\beta = \max \left\{ f_h \frac{(a_h - c_2)^2}{4}, \frac{(f_l(a_l - c_2) + f_h(a_h - c_2))^2}{4} \right\} \quad (\text{A.13})$$

The latter term in (A.13) is optimal if

$$\left(\sqrt{f_h} - f_h \right) (a_h - c_2) < f_l (a_l - c_2)$$

which can be written as

$$\frac{a_h - c_2}{a_l - c_2} < \left(\frac{1 - f_h}{\sqrt{f_h} - f_h} \right) \quad (\text{A.14})$$

The ratio of profit to optimal profit is then

$$\frac{U_\beta}{U_o} = \frac{\frac{(f_l(a_l - c_2) + f_h(a_h - c_2))^2}{4}}{\frac{f_l}{4}(a_l - c_2)^2 + \frac{f_h}{4}(a_h - c_2)^2}$$

which can be written as

$$\frac{U_\beta}{U_o} = \frac{\left(f_l + f_h \left(\frac{a_h - c_2}{a_l - c_2} \right) \right)^2}{f_l + f_h \left(\frac{a_h - c_2}{a_l - c_2} \right)^2}$$

which is decreasing in

$$\frac{a_h - c_2}{a_l - c_2}$$

Thus, the profit ratio is minimized when (A.14), which after simplification yields the bound (A.12).

If the first term in (A.13) is optimal then the ratio is

$$\frac{U_\beta}{U_o} = \frac{f_h \frac{(a_h - c_2)^2}{4}}{\frac{f_l}{4}(a_l - c_2)^2 + \frac{f_h}{4}(a_h - c_2)^2}$$

which can be written as

$$\frac{U_\beta}{U_o} = \frac{f_h}{f_l \left(\frac{a_l - c_2}{a_h - c_2} \right)^2 + f_h}$$

which is increasing in

$$\frac{a_h - c_2}{a_l - c_2}$$

meaning that the ratio is again minimized at the break point (A.14). Hence, the lower bound is again (A.12).

A.2.2. Membership Fee Contract

Suppose the platform charges a membership fee, f , to all workers that join the platform and leaves all other revenues to providers, i.e. $w_j = p_j$. The platform's profit is

$$U = fN$$

and providers expect to earn

$$\Pi = \sum_{j \in \{l, h\}} \int_0^{\phi_j p_j} G(c) dc f_j - f.$$

The platform solves

$$\begin{aligned} \max_{f, N, p_l, p_h} U \\ s.t. \Pi \geq c_1. \end{aligned}$$

To identify the solution to the platform's problem, first notice that U is increasing in f while Π is decreasing in f . For given recruitment and prices, that platform selects the membership fee that sets $\Pi = c_1$, i.e.

$$f^* = \sum_{j \in \{l, h\}} \int_0^{\phi_j p_j} G(c) dc f_j - c_1.$$

Using this membership fee, the platform's profit is

$$U(f^*) = N \left(\sum_{j \in \{l, h\}} \int_0^{\phi_j p_j} G(c) dc f_j - c_1 \right).$$

To find the profit maximizing prices, we may consider the profit earned in each demand

state separately:

$$U(f^*) = \sum_{j \in \{l, h\}} U_j - c_1$$

where

$$U_j = N \int_0^{\phi_j p_j} G(c) dc f_j$$

and

$$\begin{aligned} \frac{\partial U_j}{\partial p_j} &= NG(\phi_j p_j) \left(\phi_j + p_j \frac{\partial \phi_j}{\partial p_j} \right) \\ &= \frac{NG(\phi_j p_j)}{NG(\phi_j p_j) + \phi_j p_j Ng(\phi_j p_j)} (\phi_j NG(\phi_j p_j) - p_j) \\ &= \frac{NG(\phi_j p_j)}{NG(\phi_j p_j) + \phi_j p_j Ng(\phi_j p_j)} (a_j - 2p_j) \end{aligned}$$

where the last equivalence follows from the definition of ϕ_j :

$$\phi_j = (a_j - p_j)/(NG(\phi_j p_j)).$$

It follows that, for a given level of recruitment, U_j is quasiconcave in p_j with maximizer $p_j^* = a_j/2$.

It remains to optimize over the recruitment level, N . Define $\tilde{\phi}_j = \phi_j(p_j = a_j/2) = a_j/(2NG(\tilde{\phi}_j a_j/2))$. The platform's profit as a function of recruitment is

$$U(N) = \begin{cases} N(\sum_{j \in \{l, h\}} \int_0^{a_j/2} G(c) dc f_j - c_1), & N < a_l/2G(a_l/2) \\ N(\int_0^{a_l \tilde{\phi}_l/2} G(c) dc f_l + \int_0^{a_h/2} G(c) dc f_h - c_1), & a_l/2G(a_l/2) \leq N \leq a_h/2G(a_h/2) \\ N(\sum_{j \in \{l, h\}} \int_0^{a_j \tilde{\phi}_j/2} G(c) dc f_j - c_1), & a_h/2G(a_h/2) < N \end{cases}$$

The optimal N does not fall into the first regime because $U(N)$ is strictly increasing in that domain. The optimal value of N may be revealed via a search over its feasible range.

To derive an upper bound on N , recall that with the commission model and a give com-

mission, β , optimal wages were described by $w_j^* = \max\{\hat{w}_j, \beta a_j/2\}$. Define $\phi_j(w_j)$ as a function of wages:

$$\phi_j(w_j) = (a_j - w_j/\beta)/(NG(\phi_j(w_j)w_j)).$$

Optimal recruitment with the commission contract, N_β , satisfies

$$\begin{aligned} 0 &= \sum_{j \in \{l, h\}} \int_0^{w_j^* \phi_j(w_j^*)} G(c) dc f_j - c_1 \\ &= \sum_{j \in \{l, h\}} \int_0^{\phi_j(\beta a_j/2) \beta a_j/2 * \mathbb{1}\{\beta a_j/2 > \hat{w}_j\} + \hat{w}_j \mathbb{1}\{\hat{w}_j > \beta a_j/2\}} G(c) dc f_j - c_1 \end{aligned}$$

For $\beta = 1$, N_β :

$$\begin{aligned} 0 &= \sum_{j \in \{l, h\}} \int_0^{\phi_j(a_j/2) a_j/2 * \mathbb{1}\{a_j/2 > \hat{w}_j\} + \hat{w}_j \mathbb{1}\{\hat{w}_j > a_j/2\}} G(c) dc f_j - c_1 \\ &\geq \sum_{j \in \{l, h\}} \int_0^{a_j \tilde{\phi}_j/2} G(c) dc f_j - c_1 \end{aligned}$$

Notice that the last line is the second term of $U(N : a_h/(2G(a_h/2) < N))$, i.e.

$$U(N : a_h/(2G(a_h/2) < N)) = N \left(\sum_{j \in \{l, h\}} \int_0^{a_j \tilde{\phi}_j/2} G(c) dc f_j - c_1 \right),$$

so $U(N_\beta) \leq 0$. Consequently, N_β for $\beta = 1$ is an upper bound for recruitment in the membership fee model. It follows that any valid upper bound on N_β is also a valid upper bound for recruitment in the membership fee model when evaluated at $\beta = 1$.

A.2.3. Heterogeneous c_1

Consider the scenario in which the platform faces a population of providers with heterogeneous joining cost, c_1 . Specifically, the platform faces a finite population, M , with joining cost c_L , while the rest of the population has joining cost c_H . The platform recruits providers in order of joining cost, so the joining cost the platform faces is an increasing function of

recruitment, defined as $C(N) = c_L \mathbb{1}\{N \leq M\} + c_H \mathbb{1}\{M < N\}$. Define w'_k , w''_k , and $\bar{\phi}_k$ as

$$\int_0^{w'_k} G(c)dc = c_k$$

$$\int_0^{w''_k} G(c)dc f_h = c_k$$

$$\int_0^{\bar{\phi}_k w} G(c)dc f_l + \int_0^w G(c)dc f_h = c_k.$$

Fixed Contract

As in the proof of Theorem 1, the optimal fixed contract must have $a_l - p \leq NG(w) \leq a_h - p$, and never has $NG(w) = a_h - p$ while $C(N) < \Pi$. The platform maximizes

$$U = (p - w)((a_l - p)f_l + NG(w)f_h)$$

subject to $0 \leq \phi_l \leq 1$ and $C(N) \leq \Pi$, where

$$\Pi = \int_0^{\phi_l w} G(c)dc f_l + \int_0^w G(c)dc f_h.$$

It is never optimal for $\phi_l = 1$ and $C(N) < \Pi$. To explain, suppose $\phi_l = 1$ while $C(N) < \Pi$. Then U is strictly decreasing in w while Π is strictly increasing in w . It is therefore not optimal for the constraint $C(N) \leq \Pi$ not to bind. Consequently, at optimal there exists optimal recruitment $N^* : \Pi = C(N^*)$. Let p' be the smallest price for which the marginal recruited provider has joining cost c_H for a given wage. Then optimal recruitment may be expressed as a function of price:

$$N^* = \begin{cases} G(w)/(\bar{\phi}_L G(\bar{\phi}_L w)), & p' < p \\ G(w)/(\bar{\phi}_H G(\bar{\phi}_H w)), & p \leq p' \end{cases}$$

Platform profit as a function of price is

$$U = \begin{cases} (p - w)(a_l - p)(f_l + f_h G(w)/(\bar{\phi}_L G(\bar{\phi}_L w))), & p' < p \\ (p - w)(a_l - p)(f_l + f_h G(w)/(\bar{\phi}_H G(\bar{\phi}_H w))), & p \leq p'. \end{cases}$$

The unconstrained optimal price, $(a_l + w)/2$, is unaffected by the joining cost of the marginal provider. However, price must ensure that $N^*G(w) \leq a_h - p$. The constraint imposed on price is a function of w . If w such that $p' < (a_l + w)/2$, then $p^* = \max((a_l + w)/2, (G(w)a_l - \bar{\phi}_L G(\bar{\phi}_L w)a_h)/(G(w) - \bar{\phi}_L G(\bar{\phi}_L w)))$, otherwise $p^* = \max((a_l + w)/2, (G(w)a_l - \bar{\phi}_H G(\bar{\phi}_H w)a_h)/(G(w) - \bar{\phi}_H G(\bar{\phi}_H w)))$. A search over $w \in [w'_L, \min(a_l, w''_L)] \cup [w'_H, \min(a_l, w''_H)]$ yields the optimal contract.

Dynamic Wage Contract

There are three cases to consider. First, the contract described in Theorem 2 with $c_1 = c_L$ has $N < M$. Second, Theorem 2 with $c_1 = c_L$ has $M < N$ and the platform chooses $M < N$. Third, Theorem 2 with $c_1 = c_L$ has $M < N$ and the platform chooses $N \leq M$. In the first two cases Theorem 2 characterizes the optimal contract where Theorem 2 is evaluated with $c_1 = c_L$ and $c_1 = c_H$ respectively. We show that the third case simply requires an additional constraint on the optimal price to account for the new constraint, $N \leq M$.

To characterize the dynamic wage contract in the third case, we consider without loss of generality the problem of a per participation wage instead of a per service wage. Platform profit in demand state j is

$$U_j = \begin{cases} (p - w_j)NG(w_j), & NG(w_j) \leq a_j - p \\ p(a_j - p) - w_j NG(w_j), & a_j - p < NG(w_j) \end{cases}$$

and provider profit is

$$\Pi = \sum_{j \in \{l, h\}} \int_0^{w_j} G(c) dc f_j.$$

Clearly the optimal contract must have $\min_j \{(a_j - p)/G(w_j)\} \leq N \leq \max_j \{(a_j - p)/G(w_j)\}$. Otherwise, platform profit is either monotonically increasing or decreasing in N while Π is independent of N . To see that it is never the case that $M < \min_j \{(a_j - p)/G(w_j)\}$, notice that if $M < \min_j \{(a_j - p)/G(w_j)\}$ platform profit would be monotonically increasing in p , a contradiction.

Optimal platform profit is linear in N :

$$U = \begin{cases} NG(w_h)(p - w_h)f_h + f_l p(a_l - p) - w_l NG(w_l)f_l, & (a_l - p)/G(w_l) < (a_h - p)/G(w_h) \\ NG(w_l)(p - w_l)f_l + f_h p(a_h - p) - w_h NG(w_h)f_h, & (a_h - p)/G(w_h) \leq (a_l - p)/G(w_l) \end{cases}$$

Suppose $(a_l - p)/G(w_l) < (a_h - p)/G(w_h)$. Then there are three cases. If $w_l G(w_l)f_l < G(w_h)(p - w_h)f_h$ then U is increasing in N . Consequently recruitment is M if $M < (a_h - p)/G(w_h)$, and $(a_h - p)/G(w_h)$ otherwise. In either case U is decreasing in w_l meaning either $(a_l - p)/G(w_l) = (a_h - p)/G(w_h)$ or $\Pi = c_L$. To see that it is never optimal to have $(a_l - p)/G(w_l) = (a_h - p)/G(w_h)$ and $\Pi > c_L$ consider the following. In the first case this implies that $M < (a_l - p)/G(w_l)$, a contradiction. In the second case, U is decreasing in w_h , which means $\Pi = c_L$, a contradiction.

If instead U is decreasing in N , then $N = (a_l - p)/G(w_l)$. In this case U is again decreasing in w_l , so the optimal low wage satisfies one of the following conditions: $w_l G(w_l)f_l = G(w_h)(p - w_h)f_h$, $(a_l - p)/G(w_l) = (a_h - p)/G(w_h)$, or $\Pi = c_L$. In the first case, the analysis in the previous paragraph applies. We now show that the optimal low wage does not satisfy the second case without satisfying the third. If the second condition is satisfied, U is decreasing in w_h , meaning the third case must be satisfied.

We conclude from the above analysis that at optimal $w_h(w_l) : \Pi = c_L$. We now show that

$N = (a_l - p)/G(w_l)$. There are two alternatives to rule out. The first is $N = (a_h - p)/G(w_h)$, in which case platform profit is decreasing in w_l , which violates $(a_l - p)/G(w_l) < (a_h - p)/G(w_h(w_l))$. The second is $N = M < (a_h - p)/G(w_h(w_l))$, in which case U is again decreasing in w_l , which contradicts $N = M < (a_h - p)/G(w_h(w_l))$.

A symmetric analysis applies to the case where $(a_h - p)/G(w_h(w_l)) < (a_l - p)/G(w_l)$. Specifically the optimal contract in this case has $w_h(w_l) : \Pi = c_L$ and $N = (a_h - p)/G(w_h(w_l))$. We now show that recruitment is $(a_l - p)/G(w_l)$. Suppose $N = (a_h - p)/G(w_h(w_l)) < (a_l - p)/G(w_l)$. Then platform profit is increasing in w_l , which violates $(a_h - p)/G(w_h(w_l)) < (a_l - p)/G(w_l)$.

It remains to solve

$$\begin{aligned} \max_{w_l, p} (a_l - p) & \left(\frac{G(w_h(w_l))}{G(w_l)} (p - w_h(w_l)) f_h + (p - w_l) f_l \right) \\ \text{s.t. } a_l - p & \leq MG(w_l) \\ (a_l - p)G(w_h(w_l)) & \leq (a_h - p)G(w_l). \end{aligned}$$

Platform profit is concave in p with maximizer

$$p^* = \frac{a_l}{2} + \frac{w_h(w_l)G(w_h(w_l))f_h + w_lG(w_l)f_l}{2(G(w_h(w_l))f_h + G(w_l)f_l)}.$$

Define p' and p'' as the lower bounds on price such that the constraints in the program above are satisfied respectively.

$$\begin{aligned} p' &= a_l - MG(w_l) \\ p'' &= \frac{a_lG(w_h) - a_hG(w_l)}{G(w_h) - G(w_l)}. \end{aligned}$$

The optimal price is then $\max\{p^*, p', p''\}$. A search over $w_l \in \{w'_L, w''_L\}$ reveals the optimal dynamic wage contract.

A.3. Dynamic Price Contract

As in the proof of Theorem 3, the optimal price schedule is $\max\{a_j - NG(w), (a_j + w)/2\}$. As before, the platform does not choose a contract with $p_j = (a_j + w)/2 \forall j$ or $p_j = a_j - NG(w) \forall j$. Platform profit is

$$U = \frac{(a_l - w)^2}{4} f_l + (a_h - NG(w) - w) NG(w) f_h$$

and provider expected profit is

$$\Pi = \int_0^{\phi_l w} G(c) dc f_l + \int_0^w G(c) dc f_h$$

Platform profit is increasing in N while $\Pi - C(N)$ is decreasing in N , so recruitment is the largest N satisfying $\Pi \geq C(N)$. A search over $w \in [w'_L, \min(a_l, w''_L)] \cup [w'_H, \min(a_l, w''_H)]$ yields the optimal contract.

Fixed Commission Contract

The optimal wage schedule is defined as in Theorem 4. Platform profit continues to be weakly increasing in N and provider expected profit is monotonically decreasing in N . Consequently, optimal recruitment ensures that the profit constraint binds, and a search over the commission $\in [0, 1]$ identifies the optimal contract.

Optimal Contract

There are three cases to consider. First, the optimal contract described in Theorem 5 with $c_1 = c_L$ has $N < M$. Second, Theorem 5 with $c_1 = c_L$ has $M < N$ and the platform chooses $M < N$. Third, Theorem 5 with $c_1 = c_L$ has $M < N$ and the platform chooses $N \leq M$. In the first two cases Theorem 5 characterizes the optimal contract where Theorem 5 is evaluated with $c_1 = c_L$ and $c_1 = c_H$ respectively. To characterize the optimal contract in

the third case we will show that, at optimal $\Pi = c_L$ and $a_j - p_j = NG(w_j) \quad \forall j$.

Platform profit in demand state j is

$$U_j = \begin{cases} (p_j - w_j)NG(w_j), & p_j < a_j - NG(w_j) \\ (p_j - w_j)(a_j - p_j), & a_j - NG(w_j) \leq p_j \end{cases}$$

and provider profit in demand state j is

$$\Pi_j = \begin{cases} \int_0^{w_j} G(c)dc, & p_j < a_j - NG(w_j) \\ \int_0^{\phi_j w_j} G(c)dc, & a_j - NG(w_j) \leq p_j. \end{cases}$$

Suppose the platform maximizes $U = U_l f_l + U_h f_h$ s.t. $N \leq M$ and $\Pi \geq c_1$. Platform profit is monotonically increasing in p_j for small p_j (i.e. in the first regime) and Π is independent of p_j so the platform must pick $a_j - NG(w_j) \leq p_j$ at optimal. In this regime U is decreasing in w_j , while Π is increasing in w_j , so it must be that $\Pi = c_L$ at optimal. Furthermore, U is independent of N while Π is decreasing in N in this regime. Consequently the platform can lower N to decrease w_j until either $a_j - NG(w_j) = p_j$ or $\Pi = c_L$. If w_j constrained by $a_j - NG(w_j) \leq p_j$ then the platform can further decrease N until $w_j : \Pi = c_L$. The platform is then free to decrease w_{-j} until $a_{-j} - NG(w_{-j}) = p_{-j}$. So at optimal it must be the demand and participation match in all demand states.

It remains to solve

$$\begin{aligned} & \max_{N, w_l, w_h} \sum_{j \in \{l, h\}} (a_j - w_j - NG(w_j))NG(w_j)f_j \\ & \text{s.t. } w_h : \Pi = c_L \\ & N \leq M \end{aligned}$$

Platform profit is concave in N with maximizer

$$N' = \frac{\sum_{j \in \{l, h\}} (a_j - w_j) G(w_j) f_j}{2 \sum_{j \in \{l, h\}} G^2(w_j) f_j}.$$

The platform chooses as its recruitment $\min(N', M)$. A search over $w_l \in [w'_L, w''_L]$ reveals the optimal contract.

A.4. Proofs of Chapter 3 Theorems

Definitions:

$$\hat{p}(r, \beta) : \bar{F}(\hat{p}) = rG(\beta\hat{p})$$

$$p^* = \bar{F}(p^*)/f(p^*)$$

$$r_j = s_j/d_j$$

Furthermore, denote by $\pi_{m,j}(\beta)$ the profit earned in market m (e.g. gig-economy or two-sided market) in state j with commission β . The dependence of this quantity on β is replaced with a superscript $i \in \{c, o\}$ when evaluated at the best commission for a given model (e.g. fixed commission or optimal), i.e. $\pi_{g,l}^c$ is the profit earned with the optimal fixed commission in the gig-economy in the low state. Similarly, the total profit earned with commission model i in market m evaluated at the best commission for that model is denoted by π_m^i .

Proof of Proposition 3

Without the fixed commission constraint, the platform maximizes its profit independently in each state of the world. In state j , the platform solves:

$$\max_{p_j, \beta_j} (1 - \beta_j) p_j \min\{D_j, S_j\}.$$

Denote profit in state j as π_j . Then

$$\frac{\partial \pi_j}{\partial p_j} = \begin{cases} (1 - \beta_j)d_j (\bar{F}(p_j) - f(p_j)p_j), & D_j \leq S_j \\ (1 - \beta_j)S_j, & S_j < D_j. \end{cases}$$

Profit is increasing in p_j if $S_j < D_j$, so it must be that p_j is large enough so that $D_j \leq S_j$.

In that regime,

$$\frac{\partial \pi_j}{\partial \beta_j} = -p_j D_j$$

so profit is decreasing in β_j . Consequently, it must be that $D_j = S_j$ at optimal, i.e. $p_j = \hat{p}(r_j, \beta)$.

Differentiating with respect to β_j yields

$$\begin{aligned} \frac{\partial \pi_j}{\partial \beta_j} &= -p_j D_j + (1 - \beta_j)d_j(\bar{F}(p_j) - p_j f(p_j)) \frac{\partial p_j}{\partial \beta_j} \\ &= \frac{d_j p_j f(p_j) g(\beta_j p_j)}{d_j f(p_j) + s_j \beta_j g(\beta_j p_j)} \left(-d_j \frac{\bar{F}(p_j)}{g(\beta_j p_j)} - s_j \frac{\bar{F}(p_j)}{f(p_j)} + s_j (1 - \beta_j) p_j \right) \\ &= \frac{d_j p_j f(p_j) s_j g(\beta_j p_j)}{d_j f(p_j) + s_j \beta_j g(\beta_j p_j)} \left(-\frac{G(\beta_j p_j)}{g(\beta_j p_j)} - \frac{\bar{F}(p_j)}{f(p_j)} + (1 - \beta_j) p_j \right) \end{aligned}$$

where the second line follows from

$$\frac{\partial \hat{p}}{\partial \beta}(\beta_j) = -\frac{s_j p_j g(\beta_j p_j)}{d_j f(p_j) + \beta_j s_j g(\beta_j p_j)}$$

and the third line follows from the fact that $D_j = S_j$ at p_j . From the assumption that f and g are log concave, it follows that $g(x)/G(x)$ is decreasing in x and $f(x)/\bar{F}(x)$ is increasing in x . Because p_j is decreasing in β_j and

$$\frac{\partial p_j \beta}{\partial \beta}(\beta_j) = p_j \left(1 - \frac{s_j \beta_j g(\beta_j p_j)}{s_j \beta_j g(\beta_j p_j) + d_j f(p_j)} \right) \geq 0$$

it follows that at optimal

$$-\frac{G(\beta_j p_j)}{g(\beta_j p_j)} - \frac{\bar{F}(p_j)}{f(p_j)} + (1 - \beta_j)p_j = 0 \quad (\text{A.15})$$

because the left hand side above is decreasing in β_j and so π_j is quasiconcave in β_j .

Proof of Proposition 4

Suppose $S_j < D_j$. Then profit in state j is $(1 - \beta)p_j S_j$, which is increasing in p_j (and profit in other states is independent of p_j). So at optimal it must be that in each state $D_j \leq S_j$ and price p_j is uniquely defined by $\max\{\hat{p}_j, p^*\}$, where $\hat{p}_j : D_j = S_j$ and $p^* = \bar{F}(p^*)/f(p^*)$. Notice that $\hat{p}(r_l, \beta) \leq \hat{p}(r_h, \beta)$ and $\hat{p}(r, \beta)$ is decreasing in β .

In this regime, profit is

$$\pi = \sum_j (1 - \beta)p_j D_j$$

which is decreasing in β , so it must be that $D_h = S_h$.

I now show that the fixed commission must belong to the interval $[\min_j\{\beta_j^o\}, \max_j\{\beta_j^o\}]$.

There are two cases to consider. If $p_l = \hat{p}(r_l, \beta)$, then β^c solves

$$\frac{\partial \pi}{\partial \beta} = \frac{\partial \pi_l}{\partial \beta} + \frac{\partial \pi_h}{\partial \beta} = 0 \quad (\text{A.16})$$

where

$$\frac{\partial \pi_j}{\partial \beta} = -p_j d_j \bar{F}(p_j) + (1 - \beta) d_j (\bar{F}(p_j) - p_j f(p_j)) \frac{\partial p_j}{\partial \beta}.$$

By definition of β_j^o , $\partial \pi_j / \partial \beta < 0 \forall \beta > \max_j\{\beta_j^o\}$ and $\partial \pi_j / \partial \beta > 0 \forall \beta < \min_j\{\beta_j^o\}$. Hence the solution to (A.16) must lie in the interval $[\min_j\{\beta_j^o\}, \max_j\{\beta_j^o\}]$.

If instead $p_l = p^*$, then β^c solves

$$\frac{\partial \pi}{\partial \beta} = -p^* \bar{F}(p^*) + \frac{\partial \pi_h}{\partial \beta} = 0. \quad (\text{A.17})$$

The condition $p_l = p^*$ implies that at optimal $\hat{p}(r_l, \beta^c) < p^*$ which, from the definition of β_l^o implies that $\beta_l^o < \beta^c$. And again by the definition of β_h^o , (A.17) < 0 for all $\beta > \beta_h^o$. Hence the optimal fixed commission must exist in the interval $[\min_j\{\beta_j^o\}, \max_j\{\beta_j^o\}]$.

Proof of Proposition ??

Without a fixed commission, the platform maximizes its profit in each state independently of other states. Refer to profit in state j as $\pi_j = (1-\beta_j)D_jS_j = (1-\beta_j)d_j\bar{F}(p_j)s_jG(\beta_jp_jd_j\bar{F}(p_j))$.

Then

$$\begin{aligned}\frac{\partial\pi_j}{\partial p_j} &= (1-\beta_j)\left(D_jS_j + \left(p_jS_j\frac{\partial D_j}{\partial p_j} + p_jD_j\frac{\partial S_j}{\partial p_j}\right)\right) \\ &= (1-\beta)S_j\left(D_j + p\frac{\partial D_j}{\partial p_j} + \left(D_j + p\frac{\partial D_j}{\partial p_j}\right)\left(p_jD_j\beta_j\frac{g(\beta_jp_jD_j)}{G(\beta_jp_jD_j)}\right)\right) \\ &= (1-\beta)S_jd_jf(p_j)\left(\frac{\bar{F}(p_j)}{f(p_j)} - p_j\right)\left(1 + p_jD_j\beta_j\frac{g(\beta_jp_jD_j)}{G(\beta_jp_jD_j)}\right)\end{aligned}$$

where the second line follows from

$$\frac{\partial S_j}{\partial p_j} = \beta_j s_j g(\beta_j p_j D_j) \left(D_j + p \frac{\partial D_j}{\partial p_j} \right).$$

Because

$$1 + p_j D_j \beta_j \frac{g(\beta_j p_j D_j)}{G(\beta_j p_j D_j)} \geq 0$$

$\forall \beta_j p_j D_j \geq 0$ and

$$\frac{\bar{F}(p_j)}{f(p_j)} - p_j$$

is decreasing in p_j , π_j is quasiconcave in p_j with maximizer $p_j^o = p^* \doteq \bar{F}(p^*)/f(p^*)$. Notice that p_j^o is independent of the choice of β_j and also from the state.

Now consider the platform's optimal choice of β_j :

$$\begin{aligned}\frac{\partial \pi_j}{\partial \beta_j} &= -p^* D_j S_j + (1 - \beta_j) p^* D_j \frac{\partial S_j}{\partial \beta_j} \\ &= p^* D_j s_j g(\beta_j p^* D_j) \left(-\frac{G(\beta_j p^* D_j)}{g(\beta_j p^* D_j)} + (1 - \beta_j) p^* D_j \right).\end{aligned}$$

From the log-concavity of g , π_j is quasiconcave in β_j with maximizer $\beta_j^o : (1 - \beta_j^o) p^* D_j = G(\beta_j^o p^* D_j) / g(\beta_j^o p^* D_j)$.

Proof of Proposition 6

As above, the optimization of p_j is independent of the value of β in state j , and is consequently unaffected by the restriction $\beta_j = \beta \forall j$.

To see that $\beta^c \in [\min_j \{\beta_j^o\}, \max_j \{\beta_j^o\}]$, consider

$$\frac{\partial \pi}{\partial \beta} = \sum_j p^* D_j s_j g(\beta p^* D_j) \left((1 - \beta) p^* D_j - \frac{G(\beta p^* D_j)}{g(\beta p^* D_j)} \right). \quad (\text{A.18})$$

Evaluated at $\min_j \{\beta_j^o\}$, (A.18) must be positive by the definition of the β_j^o . Similarly, evaluated at $\max_j \{\beta_j^o\}$, (A.18) must be negative by the definition of the β_j^o . It follows that there is an interior solution β^c satisfying the first order condition (A.18) = 0.

Proof of Theorem 12

Consider $P_g \doteq \pi_g^c / \pi_g^o$ as a function of α . The platform's profit depends on α via its choice of commission and via d_l . So,

$$\frac{\partial \pi_g^i}{\partial \alpha} = \frac{\partial \pi_g^i}{\partial \beta} \frac{\partial \beta^i}{\partial \alpha} + \frac{1}{\alpha} \pi_{g,l}^i = \frac{1}{\alpha} \pi_{g,l}^i,$$

where $i \in \{c, o\}$ refers to the commission model employed by the platform and $\pi_{g,l}^i$ refers

to platform profit accumulated in the low state. It follows that

$$\frac{\partial P_g}{\partial \alpha} \propto \frac{1}{\alpha} (\pi_{g,h}^o \pi_{g,l}^c - \pi_{g,h}^c \pi_{g,l}^o)$$

which depends on α only through β^c . There are two cases to consider. If $\beta_l^o \leq \beta_h^o$, then $\pi_{g,l}(\beta^c)/\alpha$ decreases in β^c , while $\pi_{g,h}(\beta^c)$ increases in β^c . Furthermore, $\partial \beta^c / \partial \alpha \propto \partial \pi_{g,l}(\beta) / \partial \beta|_{\beta^c} \leq 0$. Hence, P_g is quasiconvex in α or monotonically decreasing in α . If instead $\beta_h^o < \beta_l^o$ then $\pi_{g,l}(\beta^c)/\alpha$ increases in β^c , $\pi_{g,h}(\beta^c)$ decreases in β^c , and β^c increases in α . Hence P_g is quasiconvex in α .

I now derive a lower bound for P_g . Note that

$$\begin{aligned} P_g &= \frac{\pi_{g,l}(\beta^c) + \pi_{g,h}(\beta^c)}{\pi_g^o} \\ &\geq \frac{\pi_{g,l}(\beta_j^o) + \pi_{g,h}(\beta_j^o)}{\pi_g^o} \\ &\geq \frac{\pi_j^o}{\pi_{g,l}^o + \pi_{g,h}^o}. \end{aligned}$$

It follows that $P_g \geq \max_j \{\pi_j^o / \pi_g^o\}$. Now consider how this bound moves as a function of γ and ρ .

Lemma 2. π_j^o is increasing in $r \doteq s/d$.

Proof: Differentiating yields:

$$\begin{aligned} \frac{\partial \pi_{g,j}^o}{\partial r} &= \frac{\partial \pi_{g,j}^o}{\partial \beta} \frac{\partial \beta_j^o}{\partial r} + (1 - \beta_j^o) d_j \frac{\partial}{\partial p} p \bar{F}(p)|_{p_j^o} \frac{\partial p_j^o}{\partial r} \\ &= d_j p_j^o \bar{F}(p_j^o) \frac{\partial p_j^o / \partial r}{\partial p_j^o / \partial \beta} \\ &= d_j \frac{G^2(\beta_j^o p_j^o)}{g(\beta_j^o p_j^o)} \\ &\geq 0. \end{aligned}$$

where the second line follows from the definition of β_j^o and the third line follows from the

definition of $p_{g,j}^o$. This implies that $\pi_{g,l}^o \geq \alpha\pi_{g,h}^o$.

From Lemma 2 demonstrates that $\pi_{g,l}^o$ is decreasing in γ while $\pi_{g,h}^o$ is independent of γ . It is easy to see that then $\pi_{g,l}^o/\pi_g^o$ is decreasing in γ , and $\pi_{g,h}^o/\pi_g^o$ is increasing in γ . This bound must be maximized at extreme values of γ .

Lemma 3. $\beta_h p_h^o \geq \beta_l p_l^o$.

Proof: If $\beta_h \geq \beta_l$, then because $\hat{p}(r, \beta)$ is decreasing in r and $\hat{p}\beta$ is increasing in β , it must be that $\beta_l p_l^o \leq \beta_l \hat{p}(r_h, \beta_l) \leq \beta_h p_h^o$. If alternatively $\beta_l > \beta_h$, then

$$\frac{G(\beta_h p_h^o)}{p_h^o g(\beta_h p_h^o)} + \frac{\bar{F}(p_h^o)}{p_h^o f(p_h^o)} = 1 - \beta_h > 1 - \beta_l = \frac{G(\beta_l p_l^o)}{p_l^o g(\beta_l p_l^o)} + \frac{\bar{F}(p_l^o)}{p_l^o f(p_l^o)},$$

which implies

$$\frac{G(\beta_h p_h^o)}{g(\beta_h p_h^o)} > \frac{p_h^o}{p_l^o} \left(\frac{G(\beta_l p_l^o)}{g(\beta_l p_l^o)} + \frac{\bar{F}(p_l^o)}{f(p_l^o)} \right) - \frac{\bar{F}(p_h^o)}{f(p_h^o)}.$$

To see that the right hand side above is greater than $G(\beta_l p_l^o)/g(\beta_l p_l^o)$, notice that if $\beta_h < \beta_l$, then it must be that $p_l^o < p_h^o$. If this were not the case, then $(1 - \beta_l)p_l^o \bar{F}(p_l^o) < (1 - \beta_h)p_h^o \bar{F}(p_h^o)$, which contradicts the claim proved above that π_j^o is increasing in r . Remembering the log-concavity of $f()$, it follows that $1 < p_h^o/p_l^o$ and $0 < \bar{F}(p_l^o)/f(p_l^o) - \bar{F}(p_h^o)/f(p_h^o)$. Because $g()$ is log-concave, $G(\beta_h p_h^o)/g(\beta_h p_h^o) > G(\beta_l p_l^o)/g(\beta_l p_l^o)$ implies $\beta_h p_h^o \geq \beta_l p_l^o$.

To see how the lower bound of P_g responds to changes in ρ , consider

$$\frac{\partial \pi_{g,l}^o / \pi_g^o}{\partial \rho} \propto \pi_{g,h}^o \alpha \frac{G^2(\beta_l p_l^o)}{\gamma g(\beta_l p_l^o)} - \pi_{g,l}^o \frac{G^2(\beta_h p_h^o)}{g(\beta_h p_h^o)}$$

From Lemma 3, $\pi_{g,l}^o \geq \alpha\pi_{g,h}^o$. If $p_l^o > p_h^o$ then

$$\frac{G^2(\beta_l p_l^o)}{\gamma g(\beta_l p_l^o)} = \bar{F}(p_l^o) \frac{G(\beta_l^o)}{\rho g(\beta_l^o)} < \bar{F}(p_h^o) \frac{G(\beta_h p_h^o)}{\rho g(\beta_h p_h^o)} = \frac{G^2(\beta_h p_h^o)}{g(\beta_h p_h^o)}.$$

So if $p_l^o > p_h^o$, then $\pi_{g,l}^o/\pi_g^o$ is decreasing in ρ . If instead $p_h^o \geq p_l^o$, then

$$\begin{aligned}
\frac{\partial \pi_{g,l}^o/\pi_g^o}{\partial \rho} &\propto \frac{\alpha \rho}{\gamma} (1 - \beta_h^o) p_h^o G(\beta_h^o p_h^o) \frac{G^2(\beta_l^o p_l^o)}{\gamma g(\beta_l^o p_l^o)} - \alpha \frac{\rho}{\gamma} (1 - \beta_l^o) p_l^o G(\beta_l^o p_l^o) \frac{G^2(\beta_h^o p_h^o)}{g(\beta_h^o p_h^o)} \\
&= (1 - \beta_h^o) p_h^o \frac{G(\beta_l^o)}{g(\beta_l^o)} - (1 - \beta_l^o) p_l^o \frac{G(\beta_h^o)}{g(\beta_h^o)} \\
&= \frac{\bar{F}(p_h^o) G(\beta_l^o p_l^o)}{f(p_h^o) g(\beta_l^o p_h^o)} - \frac{\bar{F}(p_l^o) G(\beta_h^o p_h^o)}{f(p_l^o) g(\beta_h^o p_h^o)} \\
&\leq 0.
\end{aligned}$$

If $p_h^o > p_l^o$ then $\pi_{g,l}^o/\pi_g^o$ is again decreasing in ρ . Notice that $\frac{\partial \pi_{g,l}^o/\pi_g^o}{\partial \rho} \propto -\frac{\partial \pi_{g,h}^o/\pi_g^o}{\partial \rho}$. It follows that $\pi_{g,h}^o/\pi_g^o$ is increasing in ρ .

Proof of Theorem 13

Because β^c is independent of ρ , clearly P_t is also independent of ρ .

As shown in Corollary 2, $P_t = 1$ if $\alpha = 1$. Furthermore, as $\alpha \rightarrow 0$, the platform faces only one state of the world, so $P_t \rightarrow 1$. It follows that P_t is minimized in α at a value in the range $(0, 1)$.

To see that P_t is quasiconvex in γ , observe that

$$\frac{\partial P_t}{\partial \gamma} \propto -\gamma \pi_{t,l}^c \pi_{t,h}^o + \gamma \pi_{t,h}^c \pi_{t,l}^o \tag{A.19}$$

where

$$\pi_{t,j}^i = (1 - \beta_j^i) d_j s_j p^* \bar{F}(p^*) G(\beta_j^i p^* d_j \bar{F}(p^*)).$$

Notice that the right side of (A.19) depends on γ only via β^c . There are two possible cases to consider. First, if $\beta_l^o \leq \beta_h^o$, then $\pi_{t,l}^c$ is decreasing in β^c , $\pi_{t,h}^c$ is increasing in β^c , and β^c is increasing in γ . It follows that (A.19) is increasing in γ . If instead, $\beta_h^o < \beta_l^o$, then $\pi_{t,l}^c$ is increasing in β^c , $\pi_{t,h}^c$ is decreasing in β^c , and β^c is decreasing in γ . It follows that (A.19) is increasing in γ . I conclude that P_t is quasiconvex in γ .

BIBLIOGRAPHY

- G. Allon, A. Bassamboo, and E. B. Çil. Large-scale service marketplaces: The role of the moderating firm. *Management Science*, 58(10):1854–1872, 2012.
- M. Armstrong. Competition in two-sided markets. *The RAND Journal of Economics*, 37(3):668–691, 2006.
- B. Ata and T. Olsen. Near-optimal dynamic lead-time quotation and scheduling under convex-concave customer delay costs. *Operations Research*, 57(3):753–768, 2009.
- V. Babich, A. N. Burnetas, and P. H. Ritchken. Competition and diversification effects in supply chains with supplier default risk. *Manufacturing & Service Operations Management*, 9:123–146, 2007.
- M. Bagnoli and T. Bergstrom. Log-concave probability and its applications. *Economic theory*, 26(2):445–469, 2005.
- J. Bai, K. C. So, C. Tang, X. M. Chen, and H. Wang. Coordinating supply and demand on an on-demand service platform with impatient customers. University of California at Irvine working paper, 2016.
- S. Banerjee, R. Johari, and C. Riquelme. Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 639–639. ACM, 2015.
- D. Barnes-Schuster, Y. Bassok, and R. Anupindi. Coordination and flexibility in supply contracts with options. *Manufacturing & Service Operations Management*, 4:171–207, 2002.
- F. Bernstein and A. Federgruen. Decentralized supply chains with competing retailers under demand uncertainty. *Management Science*, 51(1):18–29, 2005.
- K. Bimpikis, O. Candogan, and S. Daniela. Spatial pricing in ride-sharing networks. 2016.
- G. Cachon and M. Lariviere. Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Science*, 47:629–646, 2001.
- G. Cachon and F. Zhang. Procuring fast delivery: Sole sourcing with information asymmetry. *Management Science*, 52:881–896, 2006.
- G. P. Cachon and M. A. Lariviere. Supply chain coordination with revenue-sharing contracts: strengths and limitations. *Management science*, 51(1):30–44, 2005.
- G. P. Cachon, K. M. Daniels, and R. Lobel. The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, 2017.

- S. Celik and C. Maglaras. Dynamic pricing and lead-time quotation for a multiclass make-to-order queue. *Management Science*, 54(6):1132–1146, 2008.
- H.-P. Chao and R. Wilson. Priority service: Pricing, investment, and market organization. *The American Economic Review*, 77:899–916, 1987.
- H.-P. Chao, S. S. Oren, S. A. Smith, and R. Wilson. Multilevel demand subscription pricing for electric power. *Energy Economics*, pages 199–217, 1986.
- M. K. Chen and M. Sheldon. Dynamic pricing in a labor market: surge pricing and flexible work on the uber platform. University of California, Los Angeles working paper, 2015.
- Y. Chen and M. Hu. Pricing and matching with forward-looking buyers and sellers. 2016.
- M. C. Cohen, R. Lobel, and G. Perakis. The impact of demand uncertainty on consumer subsidies for green technology adoption. *forthcoming Management Science*, 2015.
- P. Cohen, R. H., and J. H. Using big data to estimate consumer surplus: the case of uber. University of Oxford working paper, 2016.
- J. Cramer and A. B. Krueger. Disruptive change in the taxi business: the case of uber. Harvard University working paper available at <http://www.nber.org/papers/w22083>, 2016.
- M. A. Crew, C. S. Fernando, and P. R. Kleindorfer. The theory of peak-load pricing: A survey. *Journal of Regulatory Economics*, 8:212–248, 1995.
- Economist. Pricing the surge. March 29, 2014. <http://www.economist.com/news/finance-and-economics/21599766-microeconomics-ubers-attempt-revolutionise-taxi-markets-pricing-surge>.
- L. Einav, C. Farronato, and J. Levin. Peer-to-peer markets. Technical report, National Bureau of Economic Research, 2015.
- G. D. Eppen and A. V. Iyer. Backup agreements in fashion buying - the value of upstream flexibility. *Management Science*, 43:1469–1484, 1997.
- H. S. Farber. Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics*, 130(4):1975–2026, 2015.
- D. Farrell and F. Greig. The online platform economy: What is the growth trajectory? March 2016. URL <https://www.jpmorganchase.com/corporate/institute/institute-insights.htm>.
- S. P. Fraiberger and A. Sundararajan. Peer-to-peer rental markets in the sharing economy. *NYU Stern School of Business Research Paper*, 2015.
- I. L. Gale and T. J. Holmes. Advance-purchase discounts and monopoly allocation of capacity. *The American Economic Review*, 83(1):135–146, 1993.

- V. Goyal, G. Iyengar, and Z. Qiu. Near-optimal execution policies for demand-response contracts in electricity markets. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 3697–3702. IEEE, 2013.
- I. Gurvich, M. Lariviere, and A. Moreno. Operations in the on-demand economy: Staffing services with self-scheduling capacity. *Available at SSRN 2336514*, 2015.
- J. Hall, C. Kendrick, and C. Nosko. The effects of ubers surge pricing: A case study. *The University of Chicago Booth School of Business*, 2015.
- J. V. Hall and A. B. Krueger. An analysis of the labor market for ubers driver-partners in the united states. *Princeton University Industrial Relations Section Working Paper*, 587, 2015a.
- J. V. Hall and A. B. Krueger. An analysis of the labor market for uber’s driver-partners in the united states. *available at <http://arks.princeton.edu/ark:/88435/dsp010z708z67d>*, 2015b.
- Y. Hong and P. A. Pavlou. Is the world truly ‘flat’? empirical evidence from online labor markets. *Arizona State University working paper*, 2014.
- M. Hu and Y. Zhou. Dynamic matching in a two-sided market. University of Toronto working paper, 2015.
- M. Hu and Y. Zhou. Price, wage and fixed commission in on-demand matching. 2017.
- E. Huet. Uber raises uberx commission to 25% in five more markets. September 2015. URL <http://www.forbes.com/sites/ellenhuet/2015/09/11/uber-raises-uberx-commission-to-25-percent-in-five-more-markets>. Last accessed May 4, 2017.
- R. Ibrahim and K. Arifoglu. Managing large service systems with self-scheduling agents. University College London working paper, 2015.
- M. Isaac and N. Singer. California says uber driver is employee, not a contractor. *New York Times*, June 17, 2015. <https://www.nytimes.com/2015/06/18/business/uber-contests-california-labor-ruling-that-says-drivers-should-be-employees.html>.
- A. V. Iyer, V. Deshpande, and Z. Wu. A postponement model for demand management. *Management Science*, 49:983–1002, 2003.
- A. Kabra, E. Belavina, and K. Girotra. Peer-to-peer marketplaces: Get em’ up and running. INSEAD working paper, 2015.
- L. F. Katz and A. B. Krueger. The rise and nature of alternative work arrangements in the united states, 1995-2005. *available at <http://krueger.princeton.edu/pages/projects>*, 2016.

- J. Kim and R. S. Randhawa. Asymptotically optimal dynamic pricing in observable queues. University of Southern California working paper, 2015.
- S.-H. Kim. Strategic reliability investments in multi-indenture supply chains. 2011.
- S.-H. Kim and S. Netessine. Collaborative cost reduction and component procurement under information asymmetry. *Management Science*, 59(1):189–206, 2013.
- S.-H. Kim, M. A. Cohen, and S. Netessine. Reliability or inventory? analysis of product support contracts in the defense industry. 2010.
- M. Kosoff. A new york city politician wants to ban uber’s surge pricing - but that’s a terrible idea. *Business Insider*, March 7, 2015. <http://www.businessinsider.com/banning-ubers-surge-pricing-is-a-terrible-idea-2015-2>.
- K. Kroft and D. G. Pope. Does online search crowd out traditional search and improve matching efficiency? evidence from craigslist. *Journal of Labor Economics*, 32(2):259–303, 2014.
- Q. Liu and D. Zhang. Dynamic pricing competition with strategic consumers under vertical product differentiation. *Management Science*, 59:84–101, 2012.
- S. Liu, K. C. So, and F. Zhang. Effect of supply reliability in a retail setting with joint marketing and inventory decisions. *Manufacturing & Service Management*, 12:19–34, 2010.
- N. G. Mankiw and M. D. Whinston. Free entry and social inefficiency. *The RAND Journal of Economics*, 17(1):48–58, 1986.
- M. G. Marchand. Pricing power supplied on an interruptible basis. *European Economic Review*, 5:263–274, 1974.
- A. Moreno and C. Terwiesch. Doing business with strangers: Reputation in online service marketplaces. *Information Systems Research*, 25(4):865–886, 2014.
- R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- S. S. Oren and J. A. Doucet. Interruption insurance for generation and distribution of electric power. *Journal of Regulatory Economics*, 2:5–19, 1990.
- S. S. Oren and S. A. Smith. Design and management of curtailable electricity service to reduce annual peaks. *Operations Research*, 40:213–228, 1992.
- J.-C. Rochet and J. Tirole. Platform competition in two-sided markets. *Journal of the european economic association*, 1(4):990–1029, 2003.

- J.-C. Rochet and J. Tirole. Two-sided markets: a progress report. *The RAND Journal of Economics*, 37(3):645–667, 2006.
- N. Scheiber. Growth in the 'gig economy' fuels work force anxieties, July 12 2015. URL <http://www.nytimes.com/2015/07/13/business/rising-economic-insecurity-tied-to-\decades-long-trend-in-employment-practices.html>.
- J. Seaman and P. Odell. On goldstein's variance bound. *Advances in Applied Probability*, 17(3):pp. 679–681, 1985. ISSN 00018678.
- R. Seamans and F. Zhu. Responses to entry in multi-sided markets: The impact of craigslist on local newspapers. *Management Science*, 60(2):476–493, 2013.
- S. Smith. Efficient menu structure for pricing interruptible electric power service. *The Journal of Regulatory Economics*, 1:203–223, 1989.
- E. M. Snir and L. M. Hitt. Costly bidding in online markets for it services. *Management Science*, 49(11):1504–1520, 2003.
- M. Stoller. Uber's algorithmic monopoly, April 2014. URL <http://mattstoller.tumblr.com/post/82233202309/ubers-algorithmic-monopoly-we-are-not-setting>. Date last accessed: May 4, 2017.
- J. A. Taylor and J. L. Mathieu. Index policies for demand response. *IEEE Transactions on Power Systems*, 29:1287–1295, 2014.
- T. Taylor. On-demand service platforms. 2016.
- B. Tomlin. On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science*, 52:639–657, 2006.
- A. Tsay and W. S. Lovejoy. Quantity flexibility contracts and supply chain performance. *Manufacturing & Service Operations Management*, 1:89–111, 1999.
- Y. Wang, W. Gilland, and B. Tomlin. Mitigating supply risk: Dual sourcing or process improvement? *Manufacturing & Service Operations Management*, 12:489–510, 2010.
- E. G. Weyl. A price theory of multi-sided platforms. *The American Economic Review*, 100(4):1642–1672, 2010.
- O. Q. Wu and V. Babich. Unit-contingent power purchase agreement and asymmetric information about plant outage. *Manufacturing & Service Operations Management*, 14: 245–261, 2012.
- O. Q. Wu and R. Kapuscinski. Curtailing intermittent generation in electrical systems. *Manufacturing & Service Operations Management*, 15(4):578–595, 2013.

- Z. Yang, G. Aydin, V. Babich, and D. Beil. Supply disruptions, asymmetric information, and a backup production option. *Management Science*, 55:192–209, 2009.
- Z. Yang, G. Aydin, V. Babich, and D. Beil. Using a dual-sourcing option in the presence of asymmetric information about supplier reliability: Competition vs. diversification. *Manufacturing & Service Operations Management*, 14:202–217, 2012.
- C. A. Yano and H. L. Lee. Lot sizing with random yields: A review. *Operations Research*, 43:311–334, 1995.
- H. Yoganarasimhan. The value of reputation in an online freelance marketplace. *Marketing Science*, 32(6):860–891, 2013.
- G. Zervas, D. Proserpio, and J. Byers. The rise of the sharing economy: Estimating the impact of airbnb on the hotel industry. *Journal of Marketing Research*, (2013-16), 2016.