

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

G rard P. Cachon, Kaitlin M. Daniels, Ruben Lobel*

December 2, 2015; revised

June 14, 2016

Abstract

Recent platforms, like Uber and Lyft, offer service to consumers via “self-scheduling” providers who decide for themselves how often to work. These platforms may charge consumers prices and pay providers wages that both adjust based on prevailing demand conditions. For example, Uber uses a “surge pricing” policy, which pays providers a fixed commission of its dynamic price. We find that the optimal contract substantially increases the platform’s profit relative to contracts that have a fixed price or fixed wage (or both) and although surge pricing is not optimal, it generally achieves nearly the optimal profit. Despite its merits for the platform, surge pricing has been criticized in the press and has garnered the attention of regulators due to concerns for the welfare of providers and consumers. However, we find that providers and consumers are generally better off with surge pricing because providers are better utilized and consumers benefit both from lower prices during normal demand and expanded access to service during peak demand. We conclude, in contrast to popular criticism, that all stakeholders can benefit from the use of surge pricing on a platform with self-scheduling capacity.

Keywords: self-scheduling capacity, peer-to-peer markets, contract design, dynamic pricing, service operations, ride-sharing

1 Introduction

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,

*Operations, Information, and Decisions Department, Wharton School, University of Pennsylvania, email: {cachon, kaitd, rlobel}@wharton.upenn.edu. Thanks is extended to seminar participants at Harvard University, Northwestern University, University of California Los Angeles, University of Pittsburg, University of North Carolina Chapel Hill, and Washington University.

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 (2015)). 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 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 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 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 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 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), and (vii) the platform’s

price and wage can depend on the current level of demand.

There are two key features of the model that make this environment distinctive. 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 *demand rationing* can occur because too few providers may choose to work. However, it is also possible that *capacity rationing* can occur, because too many providers may choose to work leaving some underutilized. 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.

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 later 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. 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 optimal version of each of these contracts is unavailable, but we analytically determine how to determine the optimal 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.

Beyond the contractual form, our framework allows us to explore the consequence of self-scheduling. As already discussed, self-scheduling imposes a challenge for the platform because it can only regulate the level of participation through its contract rather than directly, as it would with “central-scheduling.” However, self-scheduling provides a benefit to the providers: they can work when it is most desirable for them to do so, rather than being forced to work at times the platform dictates.

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. However, again relative to the fixed contract, adding dynamic pricing and dynamic wages together generally can increase consumer surplus even though that combination also maximizes the platform’s profit. It does so when the fixed contract rations demand when demand is high. 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 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.

There is a large number of papers in operations management that study both capacity and pricing within the same model. In the price-setting newsvendor model, a single capacity and price decision is made before demand uncertainty is resolved (see Petruzzi and Dada (1999) for a comprehensive review). In a supply chain setting, “wages” are used to align firm and provider incentives, though a wage-only (traditionally referred to as price-only) contract cannot fully coordinate the system (Perakis and Roels (2007), Cachon (2003)). Additional work considers dynamic prices (e.g. van Mieghem and Dada (1999b)) or dynamic wages (e.g. Cachon and Lariviere (2001), and van Mieghem and Dada (1999a)). Those papers generally have only a single provider, they do not investigate the interaction between dynamic prices and dynamic wages, and they do not explore welfare for all stakeholders. A separate body of work studies how a firm should manipulate its short-term capacity via production given that it is restricted to charge a single price over a selling horizon (e.g. Gilbert (1999, 2000)). In these papers the firm directly controls production, so there is no consideration of wages. Several papers study competition among multiple providers and establish that competition can lead to excessive entry (e.g. Mankiw and Whinston (1986), Besanko, Doraszelski, Lu, and Satterthwaite (2010)) 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.

Dynamic prices, or prices that respond to updated information, are a central focus for the revenue management literature (see Talluri and Van Ryzin (2006) and Elmaghraby and Keskinocak (2003) for reviews). In many papers supply is assumed to be fixed, which means prices are set over a (generally) finite horizon to maximize revenue. Some papers include supply (e.g. inventory) decisions: e.g. Gallego and Van Ryzin (1994), Federgruen and Heching (1999), Thowsen (1975), and Zabel (1970). Some work incorporates strategic consumer response to pricing (e.g. Su (2007), Aviv and Pazgal (2008), and Cachon and Swinney (2009)), and a growing number of papers investigate how pricing and supply decisions affect consumer surplus: e.g. Cohen, Lobel, and Perakis (2015), Yu, Debo, and Kapuscinski (2013). Those papers assume a constant marginal cost of supply and, more importantly, the platform directly controls supply. In our model the platform lacks such control: the platform must make joining the platform attractive enough in the long run and providers decide whether they want to work or not in the short run.

A set of papers considers peak-load pricing, the practice of charging higher prices during peak periods of demand (e.g. Steiner (1957), Crew, Fernando, and Kleindorfer (1995), and 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. Celik and Maglaras (2008), Ata and Olsen (2009), and Kim and Randhawa (2015) each define a model with a platform that serves stochastically arriving price sensitive consumers, but with fixed capacity. Banerjee, Riquelme, and Johari (2015) considers the value of dynamic pricing in a model with random arrivals of consumers and providers. Providers join based on their expectation of long run earnings (as in our model) and are paid a fixed commission of the consumer price (whereas we consider other wage structures). Prices are selected in the short-term in response to the length of the queue of providers active in the market but waiting for work. 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 have predictably higher demand than others for a given price.

Peer-to-peer service platforms have attracted significant academic interest; e.g. Kabra, Belavina, and Girotra (2015), Gefen and Carmel (2008), 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 homogenous and do not build reputations.

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 (2014), 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. However, to the extent that existing firms use fixed prices and wages, we show that they would be at a substantial disadvantage relative to a platform that uses dynamic prices and wages.

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. However, we do consider rationing if supply exceeds

demand or vice-versa.

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. Consumer prices are implicitly assumed to be fixed and the arrival of potential consumers is exogenous. Rationing of demand can occur via customer abandonment. Unlike in our model, the platform can directly control the number of providers in the system and all providers must choose to work in one of the two available periods. Taylor (2015) studies a queuing system 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. His model does not consider dynamic prices or wages. Gurvich, Lariviere, and Moreno-Garcia (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.

3 Model

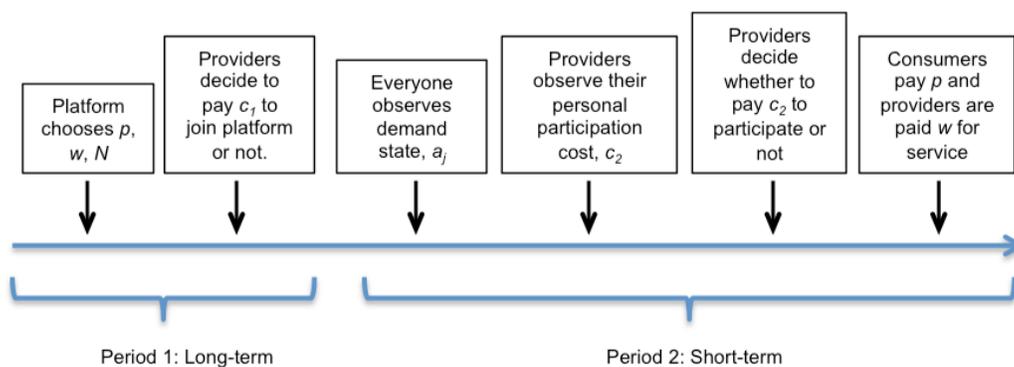


Figure 1: Timeline of events

As shown in Figure 1, 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 and without loss of generality, we collapse these decisions into a single period, which we refer to as the short-term period or the second 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. We suspect that incorporating heterogeneous opportunity costs into the model would not effect our qualitative results, but it would substantially complicate the search for optimal contracts.

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. (Recall, period 2 is really the aggregation of many decisions, and a provider’s participation cost can vary over these decisions.) 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, $G(0) = 0$, there does not exist a finite c such that $G(c) = 1$, and $g(c)/G(c)$ is decreasing, which is satisfied for log-concave distributions. Many common distributions (e.g. uniform, exponential, normal, gamma) are log-concave (see Bagnoli

and Bergstrom (2005) for additional details).

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. (Again, as period 2 represents the combination of many decisions, the demand state can vary across time.) 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 . Because feasible contracts entice all potential providers to join the platform in the first stage, N is exactly the number of recruited providers. The cap N is important - providers recognize that their earnings are decreasing in the number of providers on the platform, so the platform uses N to ensure that providers earn enough to justify the cost of joining the platform, c_1 . 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 with participation cost $\phi_j w_j$ or lower choose to participate, while providers unfortunate to have high participation costs choose not to participate. Hence, $NG(\phi_j w_j)$ providers participate. We require that providers make optimal 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, 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:

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

For a given price, wage, and recruitment level, the platform expects to earn

$$U(p, w, N) = \sum_{j \in \{l, h\}} (p_j - w_j) \phi_j NG(\phi_j w_j) f_j$$

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

$$\begin{aligned} \max_{w, p, N} & U(p, w, N) \\ \text{s.t.} & c_1 \leq \Pi(p, w, N) \end{aligned}$$

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

$$\begin{aligned} \int_0^{w'} G(c)dc &= c_1 \\ \int_0^{w''} G(c)dcf_h &= c_1 \\ \int_0^{\bar{\phi}_l w} G(c)dcf_l + \int_0^w G(c)dcf_h &= c_1 \\ \bar{c}_1 &= \sum_{j \in \{l, h\}} \int_0^{a_j} G(c)dcf_j \end{aligned}$$

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): $CS = \sum_{j \in \{l, h\}} 0.5 \min((a_j - p_j)^2, (a_j - p_j)NG(\phi_j w_j))f(a_j)$. 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 $c_1 N$.

Although our model does not include direct competition from another platform or service, it indirectly includes the effect of competition. For example, the joining cost c_1 reflects the competition for labor from other sources of employment (which may be very different than the platform's service). As the cost c_1 rises, providers earn a larger share of system surplus to reflect their greater opportunities outside of the platform. Consumers also have other options, which is why we include price sensitive demand. However, future research could explicitly model the dynamics of competition across multiple platforms.

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 optimal 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.

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 inefficiencies: 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 2. 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.

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 1 describes the optimal 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.

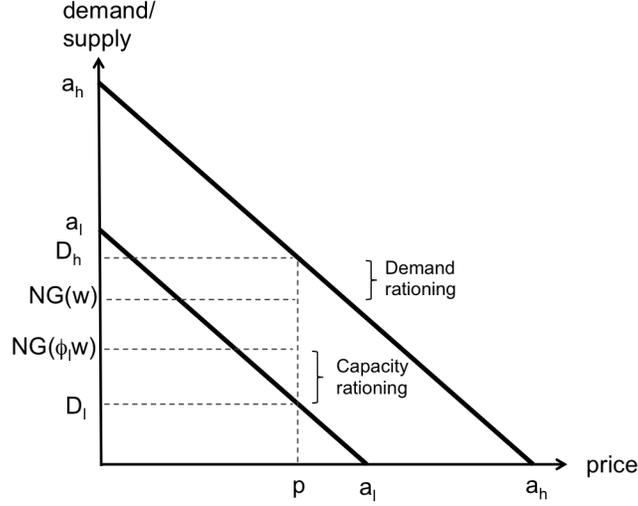


Figure 2: An example of demand and capacity rationing with a particular fixed contract.

Theorem 1. *Conditional on earning a positive profit, the optimal fixed contract has one of the following two characteristics:*

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

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

$N = (a_l - p) / (\phi_l G(\phi_l w))$, 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. $\phi_h = 1$.*

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 2, the dynamic wage contract is capable of eliminating capacity rationing in all demand states. However, the optimal dynamic wage contract may still ration demand, which is why it may not be able to earn a positive profit.

Theorem 2. *Conditional on earning a positive profit, the optimal dynamic wage contract has one of the following two characteristics:*

1. *The platform serves both demand states. In particular, $w_l \in (0, w']$,*

$$\int_0^{w_l} G(c)dc f_l + \int_0^{w_h} G(c)dc f_h = c_1$$

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

$N = (a_l - p) / G(w_l)$, 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. $\phi_h = 1$.*

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 3. *Conditional on earning a positive profit, the optimal 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 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. $c_1 = \Pi(p, w, N)$. There is no demand rationing.*

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. $\phi_h = 1$.*

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 optimal wage schedule and recruitment level satisfying the optimality conditions in the following theorem, but a line search is required to find the optimal commission.

Theorem 4. *For a given $\beta \in [\frac{w'}{a_h}, 1]$, the optimal fixed commission contract is uniquely defined, earns a positive profit for the platform and satisfies:*

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

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.

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 5. (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^o , p^o , and N^o satisfying,*

$$\begin{aligned} w_j^o &= a_j - 2N^o G(w_j^o) \\ p_j^o &= a_j - N^o G(w_j^o) \\ c_1 &= \sum_{j \in \{l, h\}} \int_0^{w_j^o} G(c) dc f_j \end{aligned}$$

(iii) *there is no capacity rationing, i.e. $\phi_l^o = \phi_h^o = 1$, 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,

$$\frac{w_j^o}{p_j^o} = \frac{w_j^o}{w_j^o + N^o G(w_j^o)}$$

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 6, 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 in order 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 6. *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) $c_1 \rightarrow 0$.*

5 Fixed Participation Cost

In this section we consider an specialized version of the main model in which, instead of heterogeneous and stochastic participation costs 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, and capacity and demand rationing are possible. The primary difference is that 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 the pool of potential providers, N , results in a selection of providers with lower average participation cost).

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.¹ For notational convenience, let $\bar{a} = f_l a_l + f_h a_h$ and $\hat{c} = c_2 + c_1/f_h$. Proofs and derivations of results are available from the authors.

The objective of this section is to use the tractability of the fixed c_2 model 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 model.

To begin the analysis, with the fixed contract it is optimal to adopt one of three versions as described in Table 1. There are two extreme solutions that serve customers in both demand states: in the “poor service” version, capacity matches low demand, thereby causing demand rationing in the high demand state; while in the “poor utilization” version, capacity matches high demand, thereby causing capacity rationing in the low demand state. In the third version only high demand is served. Thus, it is never optimal for the platform to have both demand and capacity rationing.

The optimal contract, described in Table 2, serves both demand states and sets recruitment, N , equal to high demand. To facilitate comparison with the commission contract, it is possible

¹This 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).

version	price, p	platform profit, U	consumer surplus, S
poor service: $p < a_l$, $N = a_l - p$	$\frac{a_l + c_1 + c_2}{2}$	$\frac{(a_l - c_1 - c_2)^2}{4}$	$\frac{1}{8} \left((\bar{a} - c_1 - c_2)^2 - f_h^2 (a_h - a_l)^2 \right)$
poor utilization: $p < a_l$, $N = a_h - p$	$\frac{\bar{a} + c_1 + c_2}{2}$	$\frac{1}{4} (\bar{a} - c_1 - c_2)^2 - f_l (a_h - a_l) c_1$	$\frac{1}{8} (\bar{a} - c_1 - c_2)^2 + \frac{1}{2} (f_l a_l^2 + f_h a_h^2 - \bar{a}^2)$
only high demand: $a_l < p$, $N = a_h - p$	$\frac{a_h + \hat{c}}{2}$	$\frac{1}{4} f_h (a_h - \hat{c})^2$	$\frac{1}{8} f_h (a_h - \hat{c})^2$

Table 1: Three possible versions of the optimal fixed contract in the fixed c_2 model.

	$c_1 \leq f_h (a_h - a_l)$	$f_h (a_h - a_l) < c_1$
p_l	$(a_l + c_2) / 2$	$\frac{1}{2} (a_l + c_1 + c_2 - f_h (a_h - a_l))$
p_h	$(a_h + \hat{c}) / 2$	$\frac{1}{2} (a_h + c_1 + c_2 + (1 - f_h) (a_h - a_l))$
N	$(a_h - \hat{c}) / 2$	$\frac{1}{2} (a_h - c_1 - c_2 - (1 - f_h) (a_h - a_l))$
β_l	$2c_2 / (a_l + c_2)$	$2(c_1 + c_2) / (\bar{a} + c_1 + c_2)$
β_h	$2\hat{c} / (a_l + \hat{c})$	$2(c_1 + c_2) / (\bar{a} + c_1 + c_2)$
U_o	$\frac{f_l}{4} (a_l - c_2)^2 + \frac{f_h}{4} (a_h - \hat{c})^2$	$\frac{1}{4} (f_l a_l + f_h a_h - c_1 - c_2)^2$
S_o	$\frac{1}{8} \left(f_l (a_l - c_2)^2 + f_h (a_h - \hat{c})^2 \right)$	$\frac{1}{8} (f_l a_l + f_h a_h - c_1 - c_2)^2$

Table 2: Contract terms (p_l , p_h , N , $\beta_l = w_l/p_l$, and $\beta_h = w_h/p_h$), profits (U_o) and consumer surplus (S_o) with the optimal contract for low ($c_1 \leq f_h (a_h - a_l)$) and high ($f_h (a_h - a_l) < c_1$) values of the joining cost.

to describe the platform's wage in terms of a demand-contingent commission: with low demand the platform pays a commission $\beta_l = w_l/p_l$, and with high demand it pays $\beta_h = w_h/p_h$. If $c_1 \leq f_h (a_h - a_l)$, then the amount of demand served in the low state is strictly less than the amount served in the high state. In this case the wage/commission in the low state is chosen to avoid excessive participation, leaving providers with zero earnings. (If the platform chooses a higher wage in the low state then it merely induces more participation among the providers without increasing their profit.) In the high demand state the wage/commission is selected so that all providers are willing to participate and their earnings are just sufficient to justify joining the platform. When the joining cost is higher, $f_h (a_h - a_l) < c_1$, the platform equates the amount of demand served in both states - providers are too costly to have some of them not serving customers in the low demand state. In these settings the platform uses the same commission in both demand states.

Now consider the issue of how consumer surplus compares between the fixed and optimal contracts. It can be shown analytically that whether the optimal contract increases or decreases

	high demand exceeds low demand	high demand equals low demand	only high demand
p_l	c_2/β	$\frac{1}{2}(a_l + c_1 + c_2 - f_h(a_h - a_l))$	-
p_h	\hat{c}/β	$\frac{1}{2}(a_h + c_1 + c_2 + (1 - f_h)(a_h - a_l))$	$\frac{a_h + \hat{c}}{2}$
N	$a_h - p_h$	$\frac{1}{2}(a_h - c_1 - c_2 - (1 - f_h)(a_h - a_l))$	$\frac{a_h - \hat{c}}{2}$
β	$\frac{2(f_l c_2^2 + f_h \hat{c}^2)^2}{f_l c_2 a_l + f_h \hat{c} a_h + f_l c_2^2 + f_h \hat{c}^2}$	$2(c_1 + c_2) / (\bar{a} + c_1 + c_2)$	$\frac{2\hat{c}}{a_h + \hat{c}}$
U_β	$\frac{(f_l c_2(a_l - c_2) + f_h \hat{c}(a_h - \hat{c}))^2}{4(f_l c_2^2 + f_h \hat{c}^2)}$	$\frac{1}{4}(f_l a_l + f_h a_h - c_1 - c_2)^2$	$f_h \frac{(a_h - \hat{c})^2}{4}$

Table 3: Contract terms (p_l , p_h , N , $\beta = w_l/p_l = w_h/p_h$) and profits (U_β) with three potential versions of the commission contract - each of these versions can be the best solution for the platform.

consumer surplus relative to the fixed contract depends on which version of the fixed contract yields the highest profit for the platform (i.e., the version the platform would choose to implement). If the fixed contract involves demand rationing (either poor service or only high demand) then 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 the fixed contract rations capacity (the poor utilization version), then consumers are worse off with a switch to the optimal contract. The fixed contract rations capacity when capacity is relatively cheap, meaning a low joining cost, c_1 .

The second issue is how the platform's profit compare's between the commission and optimal contracts. To be specific, when does the commission contract perform poorly relative to the optimal contract and by how much. Table 3 reports parameters and profits for three possible versions of the commission contract. (A fourth possible version exists but it does not yield closed formed solutions. This is not problematic for two reasons - it is the least likely of the versions to be optimal , and it is not necessary to include in the derivation of a profit lower bound.)

Given that we are able to evaluate the entire parameter space for the fixed c_2 model, we can determine that the profit ratio U_β/U_o is minimized when either $c_2 \rightarrow 0$ or $c_1 \rightarrow 0$. When $c_2 = 0$ and $0 < c_1$, the optimal contract chooses a low wage when demand is low (merely to induce enough participation without allowing providers to earn a profit in that state) and, when demand is high, chooses a sufficiently high wage to give providers enough profit (c_1/f_h) to justify joining the platform. Thus, the optimal low demand commission, β_l , is very small, whereas the high demand commission, β_h , is large. This disparity in the two commission rates creates a challenge for the fixed commission contract, which is required to choose a single commission. With the other extreme, $c_1 = 0$ and $0 < c_2$, the optimal policy chooses $\beta_h \ll \beta_l$. The joining constraint is not relevant (c_1

is low), so the focus for the optimal contract is on the participation constraint. Because $p_l < p_h$, the optimal commission with low demand is higher than with high demand. Again, the fixed commission contract is restricted to a single commission. If that commission is set high enough to induce participation with low demand, then it gives the providers too much profit when demand is high. But if it is set to the level needed for participation with high demand, then there will be no customers served when demand is low.

It is possible to analytically derive a lower bound (which is often tight) on the profit ratio:

$$\min \{U_\beta/U_o\} = \frac{1}{2} \left(1 + \sqrt{f_h}\right)$$

Thus, 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 $\left(\frac{1}{2} \left(1 + \sqrt{1/2}\right)\right)$.

Although there are cases in which the commission contract performs poorly relative to the optimal contract, this does require rather special parameters. For example, consider 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, $U_\beta/U_o = 0.646$. (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 $U_\beta/U_o = 0.995$ 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 the analytical results from the fixed c_2 model match the numerical analysis of the main model.

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 4 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 state ranges from a low of 0.05 to the high of 0.95. (The fixed c_2 model indicates 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 4: A summary of tested parameter values. All combinations of these values constitute 14,700 scenarios.

Table 5 reports on the frequency of different versions of the fixed contract. In 10.3% of the scenarios the fixed contract is unable to earn a positive profit for the platform. (By construction, because $c_1 < \bar{c}_1$, the optimal and commission contracts earn a positive profit in all scenarios.) In the majority of scenarios (74.3%) the fixed contract serves only the high demand state - as the fixed contract does not respond to demand, it can be better to focus on only one demand state rather than trying to cater to both. (The optimal contract always serves both demand states.) When the fixed contract serves both demand states (2,253 scenarios), it does so with one of the two extreme versions identified in section 4.1. 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. No scenarios were found which have both capacity and demand rationing. (This is always true in the fixed c_2 model.)

Table 6 reports 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

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 5: Frequency of different versions of the fixed contract.

	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

Table 6: Profit performance of the four suboptimal contracts relative to the optimal contract in all 14,700 scenarios. f , w , p , β , and o to refer to the fixed, dynamic wage, dynamic price, commission and optimal contracts, respectively.

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 for these scenarios, $U_\beta/U_o = \frac{1}{2} (1 + \sqrt{0.05}) = 0.612$.

Table 7 focuses 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 (and therefore more plausible)

	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 7: Profit performance of the four suboptimal contracts relative to the optimal contract in the 2,253 scenarios in which the fixed contract serves both demand states. f , w , p , β , and o to refer to the fixed, dynamic wage, dynamic price, commission and optimal contracts, respectively.

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 (i.e., in the 12,447 scenarios not included in the table the fixed contract earns the same profit whether the low demand state exists or not). 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).

The overall conclusions from Tables 6 and 7 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.

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, 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.

To explore a bit further the finding that the commission contract generally performs very well relative to the optimal contract, Table 8 reports the correlation between the profit ratio of the

	δ/\bar{a}	f_l	σ/μ	$G(\bar{a})$	c_1/\bar{c}_1
U_β/U_o	0.13	-0.19	0.31	0.03	0.22

Table 8: Correlations between scenario parameters and the profit ratio of the commission and optimal contracts, U_β/U_o , in the 2,253 scenarios in which the fixed contract serves both demand states.

two contracts, U_β/U_o , and the other parameters. As expected, the commission contract does best when there is less variability in demand (i.e., δ is high): the commission contract is identical to the optimal contract if there is no variability in demand. Consistent with the analytical results obtained from the fixed c_2 model, the profit ratio U_β/U_o is lowest when c_1 is low and $f_l = 1 - f_h$ is high. Finally, the commission contract performs poorly when the coefficient of variation of the participation cost, c_2 , is smallest, which suggests that the lower bound for U_β/U_o derived from the fixed c_2 model (which has zero variation in c_2) should provide a suitable bound even in the main model. Furthermore, it emphasizes that the commission contract is most effective in environments in which providers really want to work on some days, don't want to work on other days, and their is no consistent pattern to those days across time.

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, with the membership contract there is not 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. (The analogous ratios for the entire sample are 0.964 and 0.531.) 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.)

Turning to consumer surplus, we use the fixed contract as the benchmark. Tables 9 and 10

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 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 all stakeholders, 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 10). 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 9). 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.

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 results 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 10) and 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	1.001	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 9: 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.

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 10: 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.

are worse off with the optimal contract if the fixed contract chooses the poor utilization version (Table 9). The fixed c_2 model analytically yields the same result.

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

Tables 9 and 10 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

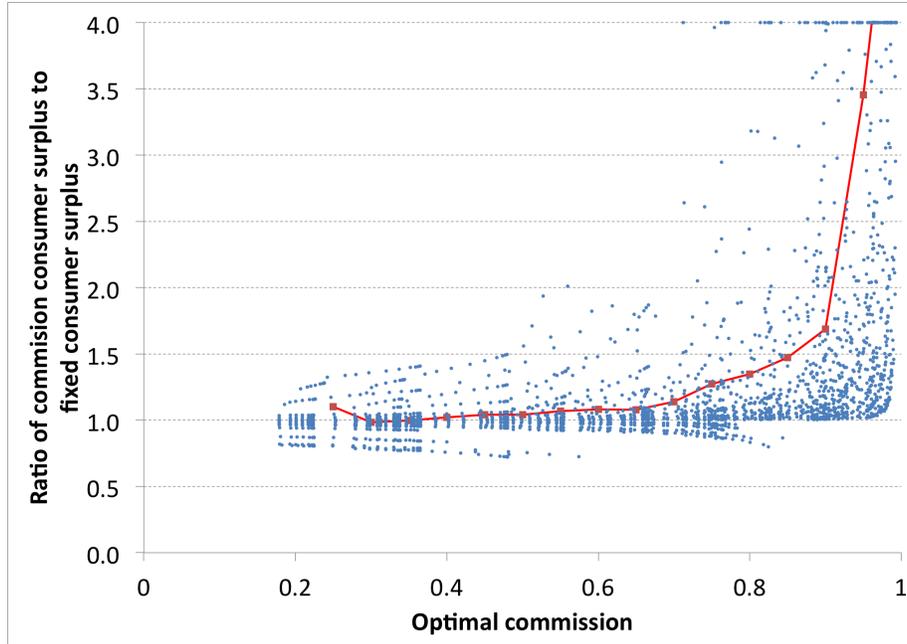


Figure 3: The ratio of consumer surplus with the commission contract to consumer surplus with the fixed contract as a function of the commission earned by providers with the commission contract in the 2,253 scenarios in which the fixed contract serves both demand states. Red squares indicate the average ratio for scenarios grouped by the commission contract commission in 0.05 intervals.

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.

7 Self-Scheduling vs. Central-Scheduling of Capacity

Self-scheduling capacity is a relatively new business practice. Traditionally, the platform decides 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 use 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.

In our model with central-scheduling the platforms selects the number of providers who participate, K_j , for each realization of available demand where K_j is bounded by recruitment, i.e. $K_j \leq N$. Participating providers earn $w_j\phi_j - c_2$ and have a K_j/N chance of being selected to participate. Recall that μ is the mean participation opportunity cost. Providers maintain their autonomy in the long-run first period. Hence, to recruit providers the platform must still compensate them for their participation costs. The platform's problem is then

$$\begin{aligned} \max_{K,w,p,N} \quad & \sum_{j \in \{l,h\}} (p_j - w_j)\phi_j K_j f_j \\ \text{s.t.} \quad & \Pi \geq c_1 \\ & N \geq K_j \forall j \end{aligned} \tag{1}$$

where the providers' expected profit from joining the platform is

$$\Pi = \sum_{j \in \{l,h\}} (w_j\phi_j - \mu) \frac{K_j}{N} f_j$$

Theorem 7. *The platform earns higher profit and providers earn higher surplus with self-scheduling than with central-scheduling of capacity.*

The platform and providers are worse off with central-scheduling because exerting direct control over provider participation increases the average provider's opportunity cost, making recruitment more difficult. The downside of self-scheduling for the platform is that it may lead to demand or capacity rationing. However the combination of dynamic wages and dynamic prices in the optimal contract eliminates both of those inefficiencies. Hence, the problem of self-scheduling can be fixed via the contractual form, but the downside of central-scheduling cannot be corrected.

This result stands in contrast to Gurvich, Lariviere, and Moreno-Garcia (2015), which shows that self-scheduling is less profitable for a platform than central-scheduling. Unlike Gurvich, Lariviere, and Moreno-Garcia (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.

Using our sample of 14,700 scenarios from the numerical study, Table 11 illustrates the magnitude of the loss that a platform incurs from central-scheduling instead of self-scheduling. On

$G(\bar{a})$	U_c/U_o					N_c/N_o				
	0.05	0.25	σ/μ 0.5	1	1.5	0.05	0.25	σ/μ 0.5	1	1.5
0.01	0.53	0.17	0.02	0.00	0.00	0.49	0.15	0.02	0.00	0.00
0.10	0.57	0.23	0.09	0.00	0.00	0.53	0.21	0.08	0.00	0.00
0.40	0.63	0.28	0.15	0.04	0.01	0.59	0.26	0.14	0.05	0.01
0.80	0.74	0.42	0.30	0.21	0.16	0.70	0.40	0.29	0.22	0.19
0.99	0.82	0.65	0.60	0.59	0.60	0.76	0.58	0.54	0.55	0.57

Table 11: The ratio of profit, U_c/U_o , and provider recruitment, N_c/N_o between the optimal contract with central-scheduling (“c” subscript) and self-scheduling (“o” subscript) in the 14,700 scenarios of the numerical study. σ/μ is the coefficient of variation of a provider’s participation cost, c_2 . $G(\bar{a})$ influences the distribution of participation cost relative to consumer willingness to pay - when $G(\bar{a})$ is small, the average participation cost, μ , is high relative to consumer willingness to pay, whereas when $G(\bar{a})$ is high, the average participation cost is low relative to consumer willingness to pay.

average, the platform’s optimal profit with central-scheduling, U_c , is only 35.7% of the optimal profit with self-scheduling providers, U_o . Providers are also worse off by a considerable amount - N_c/N_o is 33.6% on average. That said, there is considerable variation in the performance of central-scheduling. As intuition suggests, the penalty associated with central-scheduling is smaller when the providers experience less variability in their participation cost (σ/μ is small) and the average participation cost is small relative to consumer willingness to pay ($G(\bar{a})$ is large). Nevertheless, even in these scenarios, on average the platform and providers earn substantially less on average than they would with self-scheduling.

8 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. For example, demand is known to be high or low based on observable factors like the time of the week or the current weather. 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.

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. We also consider a commission contract that chooses both price and wage dynamically, but with 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. Consequently, it is probably sufficient for a platform to offer providers a fixed commission and to then dynamically only adjust the price charged to consumers.

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. We find that dynamic wages eliminate the inefficiency of capacity rationing, and thus additional supply is provided to the system. With more providers, more customers are served, leaving both providers and consumers better off. In contrast, consumers should be wary of dynamic pricing on its own. Relative to the fixed contract, adding dynamic prices while keeping the wage fixed reduces consumer welfare.

Given that dynamic pricing and dynamic wages work in opposite directions, 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.

References

- Allon, Gad, Achal Bassamboo, Eren B Çil. 2012. Large-scale service marketplaces: The role of the moderating firm. *Management Science* **58**(10) 1854–1872.
- Ata, Baris, Tava Olsen. 2009. Near-optimal dynamic lead-time quotation and scheduling under convex-concave customer delay costs. *Operations Research* **57**(3) 753–768.
- Aviv, Yossi, Amit Pazgal. 2008. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing & Service Operations Management* **10**(3) 339–359.
- Bagnoli, Mark, Ted Bergstrom. 2005. Log-concave probability and its applications. *Economic theory* **26**(2) 445–469.
- Banerjee, Siddhartha, Carlos Riquelme, Ramesh Johari. 2015. Pricing in ride-sharing platforms: A queueing-theoretic approach. *Available at SSRN 2568258* .
- Bernstein, Fernando, Awi Federgruen. 2005. Decentralized supply chains with competing retailers under demand uncertainty. *Management Science* **51**(1) 18–29.
- Besanko, David, Ulrich Doraszelski, Lauren Xiaoyuan Lu, Mark Satterthwaite. 2010. On the role of demand and strategic uncertainty in capacity investment and disinvestment dynamics. *International Journal of Industrial Organization* **28**(4) 383–389.
- Cachon, Gérard P. 2003. Supply chain coordination with contracts. *Handbooks in operations research and management science* **11** 227–339.
- Cachon, Gérard P, Martin A Lariviere. 2001. Contracting to assure supply: How to share demand forecasts in a supply chain. *Management science* **47**(5) 629–646.
- Cachon, Gérard P, Martin A Lariviere. 2005. Supply chain coordination with revenue-sharing contracts: strengths and limitations. *Management science* **51**(1) 30–44.

- Cachon, Gerard P., Robert Swinney. 2009. Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Science* **55**(3) 497–511.
- Celik, S., Costis Maglaras. 2008. Dynamic pricing and lead-time quotation for a multiclass make-to-order queue. *Management Science* **54**(6) 1132–1146.
- Chen, M. Keith, Michael Sheldon. 2015. Dynamic pricing in a labor market: surge pricing and flexible work on the uber platform. *University of Los Angeles working paper* .
- Cohen, Maxime C., Ruben Lobel, Georgia Perakis. 2015. The impact of demand uncertainty on consumer subsidies for green technology adoption. *forthcoming Management Science* .
- Cramer, Judd, Alan B. Krueger. 2016. Disruptive change in the taxi business: the case of uber. *available at <http://www.nber.org/papers/w22083>* .
- Crew, Michael A., Chitru S. Fernando, Paul R. Kleindorfer. 1995. The theory of peak-load pricing: A survey. *Journal of Regulatory Economics* **8** 212–248.
- Economist. 2014. Pricing the surge **March 29**.
- Einav, Liran, Chiara Farronato, Jonathan Levin. 2015. Peer-to-peer markets. Tech. rep., National Bureau of Economic Research.
- Elmaghraby, Wedad, Pinar Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science* **49**(10) 1287–1309.
- Farber, Henry S. 2015. 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.
- Federgruen, Awi, Aliza Heching. 1999. Combined pricing and inventory control under uncertainty. *Operations research* **47**(3) 454–475.
- Gale, Ian L, Thomas J Holmes. 1993. Advance-purchase discounts and monopoly allocation of capacity. *The American Economic Review* 135–146.
- Gallego, Guillermo, Garrett Van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science* **40**(8) 999–1020.
- Gefen, David, Erran Carmel. 2008. Is the world really flat? a look at offshoring at an online programming marketplace. *MIS quarterly* 367–384.

- Gilbert, Stephen M. 1999. Coordination of pricing and multi-period production for constant priced goods. *European Journal of Operational Research* **114**(2) 330–337.
- Gilbert, Stephen M. 2000. Coordination of pricing and multiple-period production across multiple constant priced goods. *Management Science* **46**(12) 1602–1616.
- Gurvich, Itai, Martin Lariviere, Antonio Moreno-Garcia. 2015. Operations in the on-demand economy: Staffing services with self-scheduling capacity. Tech. rep., Northwestern University, working paper.
- Hall, Jonathan V., Alan B. Krueger. 2015. An analysis of the labor market for uber’s driver-partners in the united states. *available at <http://arks.princeton.edu/ark:/88435/dsp010z708z67d>* .
- Hong, Yili, Paul A Pavlou. 2014. Is the world truly ‘flat’? empirical evidence from online labor markets. *Empirical Evidence from Online Labor Markets (October 2014)* .
- Hu, Ming, Yun Zhou. 2015. Dynamic matching in a two-sided market. *Available at SSRN 2592622* .
- Ibrahim, Rouba, Kenan Arifoglu. 2015. Managing large service systems with self-scheduling agents. *Available at SSRN 2571819* .
- Isaac, Mike, Natasha Singer. 2015. California says uber driver is employee, not a contractor. *New York Times* .
- Kabra, Ashish, Elena Belavina, Karan Girotra. 2015. Peer-to-peer marketplaces: Get em’ up and running .
- Katz, Lawrence F., Alan B. Krueger. 2016. The rise and nature of alternative work arrangements in the united states, 1995-2005. *available at <http://krueger.princeton.edu/pages/projects>* .
- Kim, Jeunghyun, Ramandeep S Randhawa. 2015. Asymptotically optimal dynamic pricing in observable queues. *Available at SSRN 2546480* .
- Kosoff, Maya. 2015. A new york city politician wants to ban uber’s surge pricing - but that’s a terrible idea. *Business Insider* .
- Kroft, Kory, Devin G Pope. 2014. Does online search crowd out traditional search and improve matching efficiency? evidence from craigslist. *Journal of Labor Economics* **32**(2) 259–303.

- Mankiw, N Gregory, Michael D Whinston. 1986. Free entry and social inefficiency. *The RAND Journal of Economics* 48–58.
- Moreno, Antonio, Christian Terwiesch. 2014. Doing business with strangers: Reputation in online service marketplaces. *Information Systems Research* **25**(4) 865–886.
- Perakis, Georgia, Guillaume Roels. 2007. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Management Science* **53**(8) 1249–1268.
- Petruzzi, Nicholas C, Maqbool Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Operations Research* **47**(2) 183–194.
- Rochet, Jean-Charles, Jean Tirole. 2006. Two-sided markets: a progress report. *The RAND Journal of Economics* **37**(3) 645–667.
- Scheiber, Noam. 2015. Growth in the 'gig economy' fuels work force anxieties. URL <http://www.nytimes.com/2015/07/13/business/rising-economic-insecurity-tied-to-decades-long-trend-in-employment-practices.html>.
- Seamans, Robert, Feng Zhu. 2013. Responses to entry in multi-sided markets: The impact of craigslist on local newspapers. *Management Science* **60**(2) 476–493.
- Snir, Eli M, Lorin M Hitt. 2003. Costly bidding in online markets for it services. *Management Science* **49**(11) 1504–1520.
- Steiner, Peter O. 1957. Peak loads and efficient pricing. *The Quarterly Journal of Economics* 585–610.
- Stoller, Matt. 2014. How uber creates an algorithmic monopoly to extract rents. URL <http://mattstoller.tumblr.com/post/82233202309/ubers-algorithmic-monopoly-we-are-not-setting>.
- Su, Xuanming. 2007. Intertemporal pricing with strategic customer behavior. *Management Science* **53**(5) 726–741.
- Talluri, Kalyan T, Garrett J Van Ryzin. 2006. *The theory and practice of revenue management*, vol. 68. Springer Science & Business Media.
- Taylor, Terry. 2015. Congested platforms. *INFORMS Annual Conference presentation* .

- Thowsen, Gunnar T. 1975. A dynamic, nonstationary inventory problem for a price/quantity setting firm. *Naval Research Logistics Quarterly* **22**(3) 461–476.
- van Mieghem, Jan, Maqbool. Dada. 1999a. Coordinating investment, production and subcontracting. *Management Science* **45**(7) 954–971.
- van Mieghem, Jan, Maqbool. Dada. 1999b. Price versus production postponement: capacity and competition. *Management Science* **45**(12) 1631–1649.
- Yoganarasimhan, Hema. 2013. The value of reputation in an online freelance marketplace. *Marketing Science* **32**(6) 860–891.
- Yu, Man, Laurens Debo, Roman Kapuscinski. 2013. Strategic waiting for consumer-generated quality information: Dynamic pricing of new experience goods. *forthcoming Management Science* .
- Zabel, Edward. 1970. Monopoly and uncertainty. *The Review of Economic Studies* 205–219.
- Zervas, Georgios, Davide Proserpio, John Byers. 2014. The rise of the sharing economy: Estimating the impact of airbnb on the hotel industry. *Boston U. School of Management Research Paper* (2013-16).

9 Appendix

9.1 Proof of Theorem 1

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 optimal contract does not exist exclusively in the first domain of the provider profit function. To explain, suppose $G(w)N < a_l - p$ were optimal. Then the platform's profit strictly increases in N while the provider's profit 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. To explain, suppose $a_h - p < G(w)N$. Utilization, 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

$$U = (p - w)(a_l - p) \left(f_l + \frac{G(w)}{\bar{\phi}_l G(\bar{\phi}_l w)} f_h \right) \quad (2)$$

and

$$\bar{p} = \frac{G(w)a_l - \bar{\phi}_l G(\bar{\phi}_l w)a_h}{(G(w) - \bar{\phi}_l G(\bar{\phi}_l w))} \leq p \quad (3)$$

As the platform profit (2) is concave in p , the optimal price, subject to the constraint (3), is $p = \max((a_l + w)/2, \bar{p})$. which satisfies the $p < a_l$ constraint as long as $w < a_l$. 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 at optimal $G(w)N \leq a_h - p$. 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 $\Pi = c_1$.

9.2 Proof of Theorem 2

See technical appendix.

9.3 Proof of Theorem 3

See technical appendix.

9.4 Proof of Theorem 5

We show that the optimal contract is equivalent to the contract that maximizes system surplus (the sum of platform and provider profit) The platform selects a price p_j and a wage w_j for each demand state a_j to maximize system 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 $p_j^o \rightarrow a_j$ and passes all profit to providers via $w_j^o \rightarrow a_j$. 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

$$\bar{c}_1 = \sum_{j \in \{l, h\}} \int_0^{a_j} G(c)dc f_j. \quad (4)$$

9.5 Proof of Theorem 4

Let wages be a fixed commission, β , of price, i.e. $w_j = \beta p_j$. As before, a provider's expected payoff for realization a_j conditional on joining is $\Pi_j = \int_0^{w_j \phi_j} G(c)dc$, where utilization in state j is implicitly defined by $\phi_j NG(\phi_j w_j) = (a_j - w_j/\beta)$. 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 $NG(\hat{w}_j) \leq a_j/2 \leq NG(\hat{w}_j)$ is an increasing function of N :

$$\frac{\partial NG(\hat{w}_j)}{\partial N} = G(\hat{w}_j) \left(1 - \frac{\beta N g(\hat{w}_j)}{\beta N g(\hat{w}_j) + 1} \right) \geq 0$$

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 $[w'/a_h, 1]$.

9.6 Proof of Theorem 6

See technical appendix.

9.7 Proof of Theorem 7

See technical appendix.