On Optimal Pricing of Services in On-demand Labor Platforms

Vijay Kamble

Department of Information and Decision Sciences
University of Illinois at Chicago
kamble@uic.edu

March 20, 2018

Abstract

I consider the optimal pricing problem faced by a freelance worker on an on-demand labor platform. Service requests arriving while the worker is busy are lost forever. Thus an hourly (or per-unit-time in general) pricing strategy appropriately captures the opportunity cost of accepting a job. With the view of improving capacity utilization in the face of arrival uncertainty, the worker may be tempted to subsidize longer jobs by offering a lower hourly rate as compared to shorter jobs, even if the customers' hourly willingness to pay is believed to be statistically identical across jobs. I show that this intuition is unfounded. If the customer arrival process is Poisson, then the optimal policy charges an identical hourly rate for all jobs irrespective of their length. This result holds even when earnings are discounted over time. I derive the optimal pricing policy in this case under standard regularity assumptions.

1 Introduction

The recent years have seen an exponential rise of the so called "gig economy", broadly consisting of online and in-person work relationships that are facilitated "on-demand" or "just-in-time" by online platforms [DS15]. While this paradigm offers new opportunities to workers for flexible employment, navigating the uncertainties associated with short-term contractual labor can be daunting. This paper is motivated by the goal of empowering the workers to take effective operational decisions in these complex and uncertain settings.

In particular, I am concerned with the revenue management challenges faced by freelance workers on platforms where the pricing decisions are individually made by the workers themselves. Examples include platforms that match semi-skilled or skilled labor to tasks of heterogeneous nature, like Upwork, Taskrabbit, or Thumbtack, but do not include ride-sharing platforms like Uber and Lyft, where the pricing is centrally regulated by the platform. These latter pricing decisions have received considerable attention from an operational perspective in recent literature [BJR15, BCD16, Tay17, CDL17, CKW17]. However, despite the presence of several blogs, articles and forums that informally and anecdotally discuss various aspects of the issue of freelancer pricing,[1] it has received almost no formal attention from the scientific community.

The central problem faced by a worker on a labor platform is that of optimally pricing her services to maximize earnings. A key feature of the on-demand service economy is that there are no

[1] A web search for "freelancer pricing strategies" returns hundreds of entries on the topic.

"per worker" queues: if a particular worker is unavailable, the customer simply chooses some other worker or leaves the platform. Hence, from the perspective of a worker, accepting a job incurs an opportunity cost of losing all the jobs that could have been accepted while the worker is busy. The longer the accepted job, the higher is this cost. Thus, a per-unit-time pricing strategy appropriately internalizes it by allowing the payment to scale with the length of the job. Such pricing is standard not only on on-demand labor platforms like Upwork, Taskrabbit, and Thumbtack (per-hour), but also on rental marketplaces like Airbnb for lodgings (per-night) or Turo/Getaround for cars (per-hour).

In this situation, the desire to improve capacity utilization in the face of arrival uncertainty may tempt a worker to offer a lower price per-unit-time for jobs that are expected to be longer as compared to the jobs that are expected to be relatively short. This may be true even if there is no significant difference in the beliefs about the customers' per-unit-time willingness to pay across different jobs. The main finding of the paper is that this intuition is incorrect: if the customers' arrival process is Poisson, and if their per-unit-time valuations are drawn from an identical distribution, the optimal pricing strategy that maximizes the long-run average earning charges a single rate for all jobs. Discounting earnings over time, i.e., giving more importance to the revenue obtained earlier, does not affect this result. Thus, any price discrimination must arise from differences in the distributions of the customers' willingness to pay.

I present this finding with some hesitancy, since once the question is appropriately formalized, the result is mathematically trivial. However, it may not be quite as obvious,[2] and could provide useful guidance to freelancers as they undertake their pricing decisions. Finally, I derive the single optimal price in this setting under standard regularity assumptions. I also present some structural insights into the optimal prices for multiple customer classes.

Related work: Pricing in service systems has been a topic of long and considerable interest in literature; see [HH03] for a survey. But the combination of heterogeneous customers and service times, and the lack of queues is, to the best of my knowledge, largely unique to the setting considered here. Three works most relevant to this setting are [ZAF06], [ZAF08] and [CSL12]. [ZAF06] considers the optimal pricing problem faced by a service facility with a single server and a queue with a finite capacity (which could be 0, as in our setting), when the customer valuations are drawn from a common distribution. Both [ZAF08] and [CSL12] analyze the case of multiple heterogeneously distributed customer classes in a similar setting. The main difference in the present work is that my focus is on the singular impact of heterogeneous job durations on optimal pricing. [ZAF06] assumes that the service time is identically distributed for all jobs, and while this assumption is not important in [ZAF08] and [CSL12], these works are focused on the case of multiple customer classes, where it is clear that price-discrimination is necessary for optimality. Moreover, even though my main result shows that job durations are irrelevant for a single customer class (conditional on adopting per-unit-time pricing), [ZAF06] derives the optimal per-job price for the case where service times are identically exponentially distributed, whereas I impose no such restrictions. Also, none of these works consider the optimization of time discounted earnings.

2 Model and Result

Consider a single worker on an on-demand service platform. Customers are of K types. Let the set of types be denoted as $\mathcal{K} \triangleq \{1, \dots, K\}$. Service requests to the worker from customers of type k

[2] For instance, price discounts for longer stays are commonly observed on rental marketplaces like Airbnb.

arrive according to a Poisson process with the rate of λ_k per hour. The durations of the jobs they bring are distributed according to G_k with mean $1/\mu_k$ hours. Let $\rho_k \triangleq \lambda_k/\mu_k$ denote the *load* of the customer channel k and let $\rho \triangleq \sum_k \rho_k$ denote the total load. Each customer i has a maximum price v_i that she is willing to pay per hour, which is drawn from a common distribution F with support on $[0, \bar{v}]$, independently across the customers. Let f denote the density function of this distribution, and let $\overline{F} = 1 - F$ denote the tail distribution function. The worker chooses an hourly rate p_k for job type k. An arriving customer's job is accepted by the worker only if she is idle and if the customer is willing to pay the hourly rate. While the worker is busy working on the job, all the arriving customers are lost forever (they are assumed to have chosen some other worker on the platform or to have left the platform altogether). While the worker is busy, she accrues a cost of c per unit time. The goal of the worker is to choose the prices p_k that maximize her long-run rate of earning.

We first derive an expression for the long-run average earning of the worker as a function of the price vector $\mathbf{p}$. Observe that since the arrival process is Poisson, the total earning until time t, denoted as $R(t)$, is a renewal reward process where each renewal cycle consists of the worker becoming idle after finishing a job, then accepting the first paying job that arrives (such jobs arrive at the rate $\overline{F}(p_k)\lambda_k$ for each type k), and then finishing the job. Let W_1 be the total earning in the first cycle and let S_1 be the length of that cycle. Then from the renewal reward theorem [Gal13], the long-run average earning of the worker as a function of $\mathbf{p}$ is given by:

$$\mathcal{R}(\mathbf{p}) = \lim_{t\to\infty} \frac{R(t)}{t} = \frac{E(W_1)}{E(S_1)} \ w.p.\,1. \tag{1}$$

Now since the arrival process is Poisson, the expected time till the first paying job arrives in a renewal cycle is $1/(\sum_{k'\in\mathcal{K}} \lambda_{k'}\overline{F}(p_{k'}))$. Also, the first paying job that arrives in a renewal cycle is of type k with probability $\lambda_k\overline{F}(p_k)/(\sum_{k'\in\mathcal{K}} \lambda_{k'}\overline{F}(p_{k'}))$. Thus we have

$$E(W_1) = \frac{\sum_{k\in\mathcal{K}} \frac{\lambda_k}{\mu_k}\overline{F}(p_k)(p_k - c)}{\sum_{k'\in\mathcal{K}} \lambda_{k'}\overline{F}(p_{k'})}$$

and

$$E(S_1) = \frac{1}{\sum_{k'\in\mathcal{K}} \lambda_{k'}\overline{F}(p_{k'})} + \frac{\sum_{k\in\mathcal{K}} \frac{\lambda_k}{\mu_k}\overline{F}(p_k)}{\sum_{k'\in\mathcal{K}} \lambda_{k'}\overline{F}(p_{k'})}.$$

We finally get

$$\mathcal{R}(\mathbf{p}) = \frac{\sum_{k\in\mathcal{K}} \rho_k(p_k - c)\overline{F}(p_k)}{1 + \sum_{k\in\mathcal{K}} \rho_k\overline{F}(p_k)}. \tag{2}$$

From this expression, it is already clear that the optimal prices do not depend on job length alone, but on the load of each customer channel. The following result is the main finding of the paper.

Theorem 1. *There exists an optimal strategy for the worker that chooses the same hourly price across all customer types.*

Proof. Recall that $\rho = \sum_{k\in\mathcal{K}} \rho_k$ and denote $\alpha_k = \rho_k/\rho$. Then for any price vector $\mathbf{p}$, we have,

$$\mathcal{R}(\mathbf{p}) = \frac{\rho\sum_{k\in\mathcal{K}} \alpha_k(p_k - c)\overline{F}(p_k)}{\sum_{k\in\mathcal{K}} \alpha_k(1 + \rho\overline{F}(p_k))} \le \rho \max_k \frac{(p_k - c)\overline{F}(p_k)}{1 + \rho\overline{F}(p_k)} \le \max_p \frac{\rho(p - c)\overline{F}(p)}{1 + \rho\overline{F}(p)}.$$

But the latter is the problem of choosing the single optimal hourly price for all jobs, given that the total load is $\rho = \sum_{k\in\mathcal{K}} \rho_k$. $\square$

2.1 Discounted earnings

At this point, one might intuitively argue that the previous finding hinges on the fact that we are optimizing the long-run average earning, and that subsidizing longer jobs may be beneficial if time is discounted, i.e., earnings earlier in the future are more important are those obtained later. Next, I show that discounting the earnings over time does not affect the result.

As before, let p_k be the hourly price for job type k. Let $X(t) \in \{0\} \cup \mathcal{K}$ be the state of the worker at time t, beginning from the state $X(0) = 0$. Here, the state 0 signifies that the worker is idle, and state $k \in \mathcal{K}$ signifies that the worker is busy with a job of type k. Then $(X(t))_{t \in \mathbb{R}_{\geq 0}}$ is a continuous time stochastic process that is càdlàg, i.e., its sample paths are right continuous with left limits. Let $R(x)$ be the earning rate in state x, where $R(0) = 0$ and $R(k) = p_k - c$ for $k \in \mathcal{K}$. Let $\gamma > 0$ be the discount factor. Then the total expected discounted reward is given by:

$$\mathcal{R}^{\gamma}(\mathbf{p}) = \mathrm{E}_{\mathbf{p}}[\int_0^{\infty} R(X(t))\mathrm{e}^{-\gamma t} dt \mid X(0) = 0]. \tag{3}$$

Let T be an exponential random variable with mean $1/\gamma$, independent of $(X(t))_{t \in \mathbb{R}_{\geq 0}}$. Consider a modified stochastic process $(\hat{X}(t))_{t \in \mathbb{R}_{\geq 0}}$, defined as $\hat{X}(t) = X(t)\mathbf{1}_{\{t<T\}} + a\mathbf{1}_{\{t \geq T\}}$ by introducing a new absorbing state $a \triangleq K + 1$ such that $R(a) = 0$. Clearly, if $(X(t))_{t \in \mathbb{R}_{\geq 0}}$ is càdlàg, then $(\hat{X}(t))_{t \in \mathbb{R}_{\geq 0}}$ is càdlàg as well. Then it is straightforward to see that $\mathcal{R}^{\gamma}(\mathbf{p})$ is the expected total earning in this modified process, i.e.,

$$\mathcal{R}^{\gamma}(\mathbf{p}) = \mathrm{E}_{\mathbf{p}}[\int_0^{\infty} R(\hat{X}(t)) dt \mid \hat{X}(0) = 0]. \tag{4}$$

Define $\beta(0) = \mathcal{R}^{\gamma}(\mathbf{p})$, and for $k \in \mathcal{K}$, define,

$$\beta(k) = \mathrm{E}_{\mathbf{p}}[\int_u^{\infty} R(\hat{X}(t)) dt \mid \hat{X}(u) = k, \hat{X}(u_-) = 0], \tag{5}$$

where $\hat{X}(u_-) = \lim_{t \uparrow u} \hat{X}(t)$. Thus, $\beta(k)$ is the expected total earning until absorption starting from the state where a job of type k has just been accepted by the worker. For each $k \in \mathcal{K}$, let $X_k \sim G_k$, and let $Y \sim \exp(1/\gamma)$, such that Y is independent of X_k. We then get the following set of first step equations for computing $\beta(0)$. First, we have,

$$\beta(0) = \frac{\sum_{k \in \mathcal{K}} \lambda_k \overline{F}(p_k)\beta(k)}{\sum_{k \in \mathcal{K}} \lambda_k \overline{F}(p_k) + \gamma}, \tag{6}$$

where $\lambda_k \overline{F}(p_k)/(\sum_{k \in \mathcal{K}} \lambda_k \overline{F}(p_k) + \gamma)$ is the probability that the process enters state k before absorption. Also, for each $k \in \mathcal{K}$, we have,

$$\beta(k) = (p_k - c)\mathrm{E}[\min(X_k, Y)] + P(X_k < Y)\beta(0). \tag{7}$$

The first term in the expression on the right is the expected earning after accepting a job until either the job is finished or the process enters the absorbing state. The expected time till either of those two events occur is $\mathrm{E}[\min(X_k, Y)]$. The second term results from the fact that the process enters state 0, i.e., the job gets finished before absorption, with probability $P(X_k < Y)$. Solving, we get,

$$\beta(0)(\sum_{k \in \mathcal{K}} \lambda_k \overline{F}(p_k) + \gamma) = \sum_{k \in \mathcal{K}} \lambda_k \overline{F}(p_k)(p_k - c)\mathrm{E}[\min(X_k, Y)] + \beta(0) \sum_{k \in \mathcal{K}} \lambda_k \overline{F}(p_k) P(X_k < Y),$$

or,

$$\beta(0)(\sum_{k\in\mathcal{K}} \lambda_k \overline{F}(p_k) P(Y \leq X_k) + \gamma) = \sum_{k\in\mathcal{K}} \lambda_k \overline{F}(p_k)(p_k - c)\mathrm{E}[\min(X_k, Y)].$$

Thus, denoting $\hat{\rho}_k = \lambda_k \mathrm{E}[\min(X_k, Y)]$, we finally have,

$$\beta(0) = \mathcal{R}^\gamma(\mathbf{p}) = \frac{\sum_k \hat{\rho}_k (p_k - c)\overline{F}(p_k)}{\gamma + \sum_{k\in\mathcal{K}} \hat{\rho}_k \overline{F}(p_k) \frac{P(Y\leq X_k)}{\mathrm{E}[\min(X_k,Y)]}}. \tag{8}$$

We now need the following fact.

Lemma 2. *Let X and Y be independent non-negative random variables such that Y is exponentially distributed with mean $1/\gamma$. Then,*

$$\frac{P(Y \leq X)}{\mathrm{E}[\min(X, Y)]} = \gamma.$$

Proof. We have,

$$\mathrm{E}[\min(X, Y) \mid X] = X \exp(-\gamma X) + (1 - \exp(-\gamma X))E(Y \mid Y \leq X). \tag{9}$$

Conditioned on the event $\{Y \leq X\}$, Y has a truncated exponential distribution, and its mean can be computed to be (see Chap. 4, Lemma 4.3 in [Oli08]),

$$E(Y \mid Y \leq X) = \frac{1}{\gamma}\left(\frac{1 - (X\gamma + 1)\exp(-X\gamma)}{1 - \exp(-X\gamma)}\right). \tag{10}$$

Thus,

$$\mathrm{E}[\min(X, Y) \mid X] = \frac{1}{\gamma}(1 - \exp(-\gamma X)). \tag{11}$$

This, coupled with the fact that $P(Y \leq X \mid X) = 1 - \exp(-\gamma X)$ implies the result. $\square$

Thus, we finally have,

$$\mathcal{R}^\gamma(\mathbf{p}) = \frac{\sum_k \hat{\rho}_k (p_k - c)\overline{F}(p_k)}{\gamma(1 + \sum_{k\in\mathcal{K}} \hat{\rho}_k \overline{F}(p_k))}. \tag{12}$$

Using the same argument as that in the proof of Theorem 1, it is straightforward to establish that there exists an optimal policy that sets the same hourly price for each customer class. This is the policy that maximizes the long-run average earning given that the total load is $\hat{\rho} = \sum_{k\in\mathcal{K}} \hat{\rho}_k$, where $\hat{\rho}_k = \lambda_k \mathrm{E}[\min(X_k, Y)]$. For instance, if the job durations are exponentially distributed, then $\hat{\rho}_k = \lambda_k/(\mu_k + \gamma)$.

2.2 The optimal strategy

When the worker simply has to set a single price p and when the total load is ρ, with some abuse of notation, her long-run average earning is,

$$\mathcal{R}(p) = \frac{\rho(p - c)\overline{F}(p)}{1 + \rho\overline{F}(p)}. \tag{13}$$

The first derivative is given by:

$$\begin{aligned}\mathcal{R}'(p) &= \frac{\Big(\rho\overline{F}(p) - (p-c)\rho f(p)\Big)(1+\rho\overline{F}(p)) + (p-c)\rho\overline{F}(p)\rho f(p)}{(1+\rho\overline{F}(p))^2} \\ &= \frac{-\rho f(p)}{(1+\rho\overline{F}(p))^2}\left(p - c - \frac{\overline{F}(p)}{f(p)} - \rho\frac{\overline{F}(p)^2}{f(p)}\right). \end{aligned} \tag{14}$$

If F has a non-decreasing hazard rate, i.e., if $f(p)/\overline{F}(p)$ is non-decreasing, then both $\overline{F}(p)/f(p)$ and $\overline{F}(p)^2/f(p)$ are non-increasing (the latter holds since $\overline{F}(p)$ is non-increasing), and hence the entire expression in the parenthesis is increasing. In this case, assuming $\bar{v} > c$, the optimal price is the unique p^* that satisfies[3]

$$p^* = c + \frac{\overline{F}(p^*)}{f(p^*)} + \rho\frac{\overline{F}(p^*)^2}{f(p^*)}. \tag{15}$$

Now, consider a firm that sells a product, with production cost of c per unit, to a population of buyers, each having a valuation independently drawn from the distribution F. It is well known (for example, see [Phi05]) that if F has a non-decreasing hazard rate, then the optimal price p^* that maximizes the firm's profit is the one that satisfies

$$p^* = c + \frac{\overline{F}(p^*)}{f(p^*)}. \tag{16}$$

The first two terms in (15) can be recognized here. The new part is the third term, which captures the externality imposed by accepting a job. If the total load ρ increases – for instance, in periods of high demand – then this externality is higher (the worker will be rejecting more jobs) and thus the optimal price increases. Another way to understand this term is to think about what happens when the worker increases the price by a small amount. The two effects that are common to product pricing are: a) everyone who can afford to pay the increased price will pay more, leading to an increase in revenue, and b) people who cannot afford to pay the increased price will not purchase the product, leading to a loss in revenue. The optimal product price balances these two effects given the distribution of valuations. But the third effect, which is unique to services, is that the customers who can afford to pay the higher price not only pay higher, *they also face less congestion and hence are rejected less frequently.* This effect is captured by the third term. The higher the load, the stronger is this effect and hence the optimal price tends to be higher.

2.3 Price differentiation with heterogeneous customers

It is now clear that any price differentiation must stem from the fact that the beliefs about the customers' willingness to pay are different across different classes. In this section, I present some structural insights into the optimal prices in this case. Let F_k be the distribution of per-unit-time valuations of customers in class $k \in \mathcal{K}$, with support on $[0, \bar{v}_k]$, and let $\bar{F}_k = 1 - F_k$ and f_k denote the corresponding tail distribution functions and densities respectively. Then it is straightforward to establish as we did earlier, that for a price vector $\mathbf{p}$, the long-run average earning is given by:

$$\mathcal{R}(\mathbf{p}) = \frac{\sum_k \rho_k (p_k - c)\overline{F}_k(p_k)}{1 + \sum_{k\in\mathcal{K}} \rho_k \overline{F}_k(p_k)}. \tag{17}$$

[3] For the case where the job durations are identically exponentially distributed, this condition can be obtained from a result in [ZAF06] (as a sub-case of the M/M/1/K model). The present derivation imposes no such restriction.

It would be convenient to define $\Phi_k(p) \triangleq \rho_k \overline{F}_k(p)$ and $\phi_k(p) \triangleq \rho_k f_k(p)$, so that we have

$$\mathcal{R}(\mathbf{p}) = \frac{\sum_k (p_k - c)\Phi_k(p_k)}{1 + \sum_{k \in \mathcal{K}} \Phi_k(p_k)}. \tag{18}$$

The partial derivative of $\mathcal{R}(\mathbf{p})$ with respect to p_1 is given by,

$$\begin{aligned}
\frac{\partial \mathcal{R}(\mathbf{p})}{\partial p_1} &= \frac{(\Phi_1(p_1) - (p_1 - c)\phi_1(p_1))(1 + \sum_{k \in \mathcal{K}} \Phi_k(p_k)) + (p_1 - c)\Phi_1(p_1)\phi_1(p_1)}{(1 + \sum_{k \in \mathcal{K}} \Phi_k(p_k))^2} \\
&\quad + \frac{\phi_1(p_1) \sum_{k \neq 1} (p_k - c)\Phi_k(p_k)}{(1 + \sum_{k \in \mathcal{K}} \Phi_k(p_k))^2} \\
&= C(\mathbf{p}) \left[p_1 - c - \frac{\sum_{k \neq 1} (p_k - c)\Phi_k(p_k)}{1 + \sum_{k \neq 1} \Phi_k(p_k)} - \frac{\Phi_1(p_1)}{\phi_1(p_1)} - \frac{\Phi_1(p_1)^2}{\phi_1(p_1)(1 + \sum_{k \neq 1} \Phi_k(p_k))} \right],
\end{aligned} \tag{19}$$

where

$$C(\mathbf{p}) = \frac{-\phi_1(p_1)(1 + \sum_{k \neq 1} \Phi_k(p_k))}{(1 + \sum_{k \in \mathcal{K}} \Phi_k(p_k))^2}. \tag{20}$$

Thus, from the first order optimality conditions (and substituting the quantities Φ_k and ϕ_k), assuming that $\mathbf{p}^*$ is in the interior of $\prod_{k \in \mathcal{K}} [0, \bar{v}_k]$, we obtain that the optimal prices $\mathbf{p}^*$ satisfy:

$$p_i^* - c = \frac{\sum_{k \neq i} (p_k^* - c)\rho_k \overline{F}_k(p_k^*)}{1 + \sum_{k \neq i} \rho_k \overline{F}_k(p_k^*)} + \frac{\overline{F}_i(p_i^*)}{f_i(p_i^*)} + \frac{\rho_i \overline{F}_i(p_i^*)^2}{f_i(p_i^*)(1 + \sum_{k \neq i} \rho_k \overline{F}_k(p_k^*))}. \tag{21}$$

The first term signifies that the earning rate while the worker is busy with a customer of class i cannot be smaller than the average earning rate obtained from the customers from all the other classes. This condition is quite natural, since if this is not feasible, it is better to just not serve that class. The second term captures the standard product pricing tradeoff: a small increase in price reduces revenue from the customers who decide to not buy at the high price and increases revenue from customers who can afford to pay the increased price. The third term is unique to the service aspect as we have discussed in the previous section: increasing the price for a customer class reduces the congestion faced by the paying customers from that class due to the reduced customer demand. But in this case, reducing "same-class" congestion increases a class's relative exposure to "cross-class" congestion, which is captured by the denominator of the third term. Thus the push towards higher prices is not as pronounced.

3 Conclusion and future directions

The flexibility and autonomy of freelance work on on-demand platforms is accompanied by the significant risk resulting from the lack of stability and benefits associated with full time employment. Effectively mitigating these risks is critical to the long-run success of the gig economy. Understanding and addressing the different operational challenges that workers face can not only directly improve short-term outcomes for the workers, but it can also help the platforms predict overall system behavior and design effective controls that improve efficiency in the long-term.

There are several interesting avenues for future research in this domain. For instance, in my analysis, I implicitly assumed an equivalence of per-unit-time prices and prices per job. Indeed, if the customers and workers both know the mean job durations and workers are assumed to honestly invest effort, then the optimal price per job is simply the product of the optimal price per-unit-time and the mean duration (assuming that the customers' willingness to pay per job is the product of the per-unit-time willingness to pay and the mean duration). But the situation could be different if there is an asymmetry of information, e.g., if the workers have better estimates of the job duration than the customers. The strategic analysis of such scenarios and the design of optimal contracts in these settings is of considerable interest to the freelancer community. From the perspective of the platform, it would also be interesting to study the price competition between workers with heterogeneous qualities and understand the behaviors that could emerge.

References

[BCD16] Kostas Bimpikis, Ozan Candogan, and Saban Daniela. Spatial pricing in ride-sharing networks. 2016.

[BJR15] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 639–639, New York, NY, USA, 2015. ACM.

[CDL17] Gerard P Cachon, Kaitlin M Daniels, and Ruben Lobel. The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, 19(3):368–384, 2017.

[CKW17] Juan Camilo Castillo, Dan Knoepfle, and Glen Weyl. Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 241–242. ACM, 2017.

[CSL12] Felipe Caro and David Simchi-Levi. Optimal static pricing for a tree network. *Annals of Operations Research*, 196(1):137–152, 2012.

[DS15] Valerio De Stefano. The rise of the just-in-time workforce: On-demand work, crowdwork, and labor protection in the gig-economy. *Comp. Lab. L. & Pol'y J.*, 37:471, 2015.

[Gal13] Robert G Gallager. *Stochastic processes: theory for applications*. Cambridge University Press, 2013.

[HH03] Refael Hassin and Moshe Haviv. *To queue or not to queue: Equilibrium behavior in queueing systems*, volume 59. Springer Science & Business Media, 2003.

[Oli08] David J Olive. Applied robust statistics. 2008.

[Phi05] Robert Lewis Phillips. *Pricing and revenue optimization*. Stanford University Press, 2005.

[Tay17] Terry Taylor. On-demand service platforms. 2017.

[ZAF06] Serhan Ziya, Hayriye Ayhan, and Robert D Foley. Optimal prices for finite capacity queueing systems. *Operations Research Letters*, 34(2):214–218, 2006.

[ZAF08] Serhan Ziya, Hayriye Ayhan, and Robert D Foley. A note on optimal pricing for finite capacity queueing systems with multiple customer classes. *Naval Research Logistics (NRL)*, 55(5):412–418, 2008.