

Revenue Management on an On-Demand Service Platform

Vijay Kamble

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

I consider the optimal hourly (or per-unit-time in general) pricing problem faced by a worker (or a service provider) on an on-demand service platform. Service requests arriving while the worker is busy are lost forever. Thus, the optimal hourly prices need to capture the average hourly opportunity costs incurred by accepting jobs. Due to potential asymmetries in these costs, price discrimination across jobs based on duration, characteristics of the arrival process, etc., may be necessary for optimality, even if the customers' hourly willingness to pay is identically distributed. I first establish that such price discrimination is not necessary if the customer arrival process is Poisson: in this case, the optimal policy charges an identical hourly rate for all jobs. This result holds even if the earnings are discounted over time. I then consider the case where the customers belong to different classes that are differentiated in their willingness to pay. I present a simple and practical iterative procedure to compute the optimal prices in this case under standard regularity assumptions on the distributions of customer valuations. I finally show that these insights continue to hold in the presence of competition between multiple quality-differentiated workers, assuming a natural customer choice model in which a customer always chooses the best available worker that she can afford.

Key words: On-Demand Services, Freelancer Pricing, Revenue Management

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 (De Stefano 2015). While this paradigm has ushered in new opportunities for flexible employment in the economy, navigating the uncertainties associated with short-term contractual work can be daunting. This paper is motivated by the goal of empowering workers and service providers on these platforms to take effective operational decisions in these complex and uncertain settings.

In particular, this paper is concerned with the revenue management challenges faced by workers on platforms where they individually make their own pricing decisions. 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 (e.g., Banerjee et al. 2015, Bimpikis et al. 2019, Taylor 2018, Cachon et al. 2017, Castillo et al. 2017). However, despite the presence of several blogs and online forums that informally and anecdotally discuss various aspects of the issue of freelancer pricing,¹ it has received little formal attention from the scientific community. To a large extent, similar revenue management challenges are faced by service providers on online rental marketplaces like Airbnb for lodgings, Turo or Getaround for cars, etc.

In this paper, I consider the central problem faced by a worker on an on-demand service platform, which is that of optimally pricing her services to maximize earnings. Because of uncertainty and heterogeneity in the job durations, workers typically use per-unit-time pricing (e.g., hourly), which allows the payment to scale linearly with the duration of a job.² Consequently, the focus is on such per-unit-time pricing strategies (I discuss the relation to per job pricing in Section 6).

A key feature of the on-demand service economy is that there are no “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 optimal price per-unit-time for a job thus needs to internalize the *per-unit-time opportunity cost*, which is the ratio of the expected earnings lost due to accepting the job and the expected duration of the job.

The *first* contribution of this paper is to highlight that, depending on system characteristics and class specific features like distributions of the job durations, particulars of the arrival processes, etc., there could be an asymmetry in the per-unit-time opportunity costs incurred by different classes of jobs. And in these cases, price discrimination across these classes may be necessary to maximize earnings even if there is no difference in the beliefs about the customers’ per-unit-time willingness to pay. The following example is a bit contrived, but it nevertheless illustrates this point. More examples of practical interest will be presented in Section 3.

Example. *Consider two types of customers: type A customers bring a job of length 1 hour and type B customers bring a job of length 2 hours. The willingness to pay per hour for each customer type is uniformly distributed in $[0, 1]$. Type A customers arrive at times $t = 0, 2, 4, 6, \dots$ and type B customers arrive at times $t = 1, 3, 5, 7, \dots$. Now, accepting a job of type A does not incur any opportunity cost since it will be finished in time to be able to accept the next job. Thus the optimal hourly price for this job can be independently computed to be $1/2$ per hour. On the other hand, accepting a customer of type B does incur an opportunity cost: the next arriving customer of type A cannot be accepted. This necessitates a reserve value on the hourly price for job type B resulting*

¹ A web search for “freelancer pricing strategies” returns hundreds of entries on the topic.

² Such pricing is standard not only on on-demand labor platforms like Upwork, Taskrabbit, and Thumbtack (hourly), but also on rental marketplaces like Airbnb (per-night) or Turo/Getaround (hourly).

from the requirement that the total earning from the job is at least as much as the optimal expected earning from the type A job that would be lost. Consequently, the optimal hourly price for job type B is higher than $1/2$.

The *second* contribution of this paper is to establish that 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. In other words, in this case, the per-unit-time opportunity costs are identical across jobs irrespective of their durations. Discounting earnings over time, i.e., giving more importance to the earnings obtained earlier, does not affect this result. Hence, any price discrimination must result from differences in the distributions of the customers' willingness to pay.

I present this particular finding with some hesitancy, since once the question is appropriately formalized, the result is mathematically trivial. However, it may not be quite as obvious, and could provide useful guidance to workers and service providers as they undertake their pricing decisions. For example, the desire to improve capacity utilization in the face of arrival uncertainty may tempt a worker to prioritize longer jobs by offering them a lower hourly price as compared to jobs that are expected to be relatively short, even if the customers' hourly willingness to pay is believed to be identically distributed across jobs.³ This temptation may be stronger if the worker is time sensitive and wishes to maximize discounted earnings. The result implies that such intuition is unfounded.

This brings us to the *third* contribution of the paper. Under the same setting of Poisson arrivals, I consider the problem of choosing the optimal set of prices for multiple customer classes that are differentiated in their per-unit-time willingness to pay. I present key structural insights into the optimal prices in this case, finally culminating in a practically attractive iterative procedure that converges to these prices under standard regularity assumptions on the distribution of customer valuations.

The analysis hinges on the observation that the per-unit-time opportunity cost of accepting a job from any customer class is in fact the optimal earning rate. If this rate is known, then the optimal pricing problem decomposes across classes into a set of independent optimal pricing decisions with a reserve price equal to the optimal earning rate. Effectively, these are basic optimal product pricing decisions. Since the optimal rate is not known, a natural iterative scheme for a worker is to start with a guess for the optimal earning rate, compute the optimal prices independently for each customer class, estimate the resulting earning rate, recompute the optimal prices, and so on. I show that under the regularity assumptions on the distribution of valuations, this procedure indeed converges to the set of optimal prices and the optimal earning rate.

³ For instance, price discounts for longer stays are commonly observed on rental marketplaces like Airbnb.

Finally, I discuss the issue of worker competition on these platforms. I show that with undifferentiated workers, an equilibrium vector of prices typically fails to exist. But for workers that are differentiated by their quality (as signaled by their ratings on these platforms, for instance), I propose a natural customer behavior model under which a price equilibrium exists, and can be efficiently computed. Under this model, that I term “Best-I-can-afford” (BICA), each arriving customer chooses the highest quality available worker that she can afford. Under this model, I show that the workers can hierarchically solve monopolistic optimal pricing problems to arrive at the price equilibrium. This in particular implies that all the structural insights into the optimal pricing problem under the monopolistic setting carry over to the competitive setting from the perspective of any single worker.

The paper is organized as follows. I discuss relevant literature in the remainder of this section. In Section 2, I present the model and the result on the optimality of a single hourly price with statistically identical customers and Poisson arrivals. In Section 3, I discuss examples of certain related settings where this prior result doesn’t hold. In Section 4, I consider the case with multiple customer classes, where I analyze the structure of the optimal prices and finally present the iterative algorithm for computing these prices. In Section 5, I consider the issue of worker competition. In Section 6, I end the paper with a discussion of possible extensions of the formulation, its applicability to practical settings, and possible future directions.

1.1. Related work

Platform pricing in two-sided markets has received considerable attention in literature. Rochet and Tirole (2003) pioneered the study of monopolistic pricing in such platforms; see also Rochet and Tirole (2006). This analysis was later generalized and extended to multi-sided platforms in Weyl (2010). Caillaud and Jullien (2003) and Armstrong (2006) additionally consider competition across platforms. These works abstract away the operational details of interactions on these platforms and exogenously specify how utilities of entities depend on the number of entities on the other side due to network effects. The focus is on traditional microeconomic questions like the impact of governance structures on the efficiency of market outcomes, whether and how prices internalize network effects, the impact of the agents’ ability to multi-home, the impact of subscription based vs. transaction based pricing, etc. Bakos and Katsamakas (2008) considers the impact of investments into modulating the strength of network effects along with pricing decisions. Later works analyze the design of pricing and fee structures in two-sided market platforms while explicitly modeling their impact on the buyer/seller incentives and the resulting competitive equilibrium prices for trading that emerge on the platform. Works in this spirit include Economides and Katsamakas (2006), Lin et al. (2011), and more recently, Birge et al. (2018). These works assume a large

market model with a continuum of entities transacting, due to which, a) price-taking behavior and competitive equilibrium emerge as natural determinants of trading prices, and b) distinction between products and services disappears.

In contrast to the platform perspective espoused in these works, the present paper focuses instead on the optimal pricing problem from the perspective of workers operating on labor platforms, while assuming the fees charged by the platforms as fixed. Moreover, in modeling competition, I assume a fixed and finite set of workers displaying fully strategic behavior in setting prices as opposed to assuming that workers are price-taking and their supply can be modulated by market prices. Because of this, the randomness inherent to the service aspect of the transactions cannot be neglected in the model, and in fact, it plays a crucial role in the analysis.

As mentioned in the introduction, the decentralized nature of these pricing decisions for trades on the platform is different from the centralized pricing seen in ride-sharing platforms and analyzed in several recent works; c.f. Banerjee et al. (2015), Bimpikis et al. (2019), Cachon et al. (2017), Castillo et al. (2017). Another important distinction in these works is the spatial nature of demand and supply, which is not modeled in the present work.

The present work is closely related to the literature on pricing in service systems; see Hassin and Haviv (2003) for a survey. Three works most relevant to the present setting are Ziya et al. (2006), Ziya et al. (2008) and Caro and Simchi-Levi (2012). Ziya et al. (2006) 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 Ziya et al. (2008) and Caro and Simchi-Levi (2012) analyze the case of multiple heterogeneously distributed customer classes in a similar setting. Although these settings are more general, the more focused analysis of the setting without queueing in the present work, motivated by on-demand platforms, results in several novel results and insights that do not hold in the general setting. In particular, the non-discrimination result with identically distributed valuations across customer classes holds only in this setting (as shown in the example in Section 3.1). Similarly, the efficient procedure for computing optimal prices with heterogeneous customers doesn't necessarily converge in the presence of queues as I show in Remark 1 at the end of Section 4. Additionally, these works do not consider the issue of price competition amongst servers, neither do they consider maximization of discounted earnings as an objective.

2. Homogeneous customers

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 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 has a maximum price v that she is willing to pay per hour, which is drawn independently from a common distribution F . Let $\bar{F} = 1 - F$ 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 $\bar{F}(p_k)\lambda_k$ 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 (Gallager 2013), the long-run average earning of the worker as a function of \mathbf{p} is given by:

$$\mathcal{R}(\mathbf{p}) = \lim_{t \rightarrow \infty} \frac{R(t)}{t} = \frac{\mathbb{E}(W_1)}{\mathbb{E}(S_1)} \text{ w.p. 1.} \quad (1)$$

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

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

and

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

We finally get

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

The following result is the first 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,

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

But the latter is the problem of choosing the single optimal hourly price for all jobs, thus proving the result. Inequality (a) relies on the following argument. Let $\Delta = \{\bar{\alpha} \in \mathbb{R}^K; \sum_{k \in \mathcal{K}} \alpha_k = 1; \alpha_k \geq 0 \forall k \in \mathcal{K}\}$. Suppose that $\mathbf{a} = (a_1, a_2, \dots, a_K)$ is a non-negative vector and $\mathbf{b} = (b_1, b_2, \dots, b_K)$ is a positive vector. Define

$$\omega^* = \max_{\bar{\alpha} \in \Delta} \frac{\sum_{k \in \mathcal{K}} \alpha_k a_k}{\sum_{k \in \mathcal{K}} \alpha_k b_k},$$

and let $\bar{\alpha}^*$ be the maximizer. Then for any $\bar{\alpha} \in \Delta$, we have

$$\omega^* \geq \frac{\sum_{k \in \mathcal{K}} \alpha_k a_k}{\sum_{k \in \mathcal{K}} \alpha_k b_k}.$$

This implies that for any $\bar{\alpha} \in \Delta$,

$$\sum_{k \in \mathcal{K}} \alpha_k (a_k - \omega^* b_k) \leq 0.$$

This inequality is an equality only when $\bar{\alpha} = \bar{\alpha}^*$. This implies that $\bar{\alpha}^*$ is the maximizer of $\sum_{k \in \mathcal{K}} \alpha_k (a_k - \omega^* b_k)$. But this implies that $\alpha_k^* > 0$ if and only if $a_k - \omega^* b_k = 0$, i.e., if $\omega^* = a_k / b_k$. This implies that $\omega^* = a_k / b_k$ for some $k \in \mathcal{K}$, which implies that $\omega^* = \max_{k \in \mathcal{K}} a_k / b_k$. \square

2.1. Discounted earnings

In this section, I show that discounting the earnings over time does not affect the previous 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:

$$\mathcal{R}^\gamma(\mathbf{p}) = \mathbb{E} \left[\int_0^\infty R(X(t)) e^{-\gamma t} dt \mid X(0) = 0 \right]. \quad (3)$$

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.,

$$\mathcal{R}^\gamma(\mathbf{p}) = \mathbb{E} \left[\int_0^\infty R(\hat{X}(t)) dt \mid \hat{X}(0) = 0 \right]. \quad (4)$$

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

$$\beta(k) = \mathbb{E}\left[\int_u^\infty R(\hat{X}(t))dt \mid \hat{X}(u) = k, \hat{X}(u_-) = 0\right], \quad (5)$$

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,

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

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

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

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 $\mathbb{E}[\min(X_k, Y)]$. 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,

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

or,

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

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

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

We now need the following fact.

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

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

Proof. We have,

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

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 Olive (2008)),

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

Thus,

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

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

Thus, we finally have,

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

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 $\hat{\rho}_k = \lambda_k \mathbb{E}[\min(X_k, Y)]$. For instance, if the job durations are exponentially distributed, then $\hat{\rho}_k = \lambda_k / (\mu_k + \gamma)$.

3. Related settings

In this section, I take a short detour and present examples to show that Poisson arrivals are not sufficient for the previous result in an arbitrary system.

3.1. Worker specific queues

Consider the case where the customers could be patient, i.e., they are willing to wait for the worker to become free before choosing someone else. This could naturally model traditional freelancing settings where the worker has developed a sustained relationship with a class of customers who prefer her services over other workers. I numerically demonstrate that in these cases, the previous result may not hold, i.e., even under Poisson arrivals, price discrimination may be necessary for optimality despite identically distributed customer preferences.

Suppose that there are two classes of customers, A and B. The hourly willingness to pay of each customer is uniformly distributed in $[0, 1]$. Customers from class A arrive at the rate of λ_A customers per hour and those from B arrive at the rate of λ_B customers per hour. The job durations of both classes are exponentially distributed: class A with mean $1/\mu_A$ hours and class B with mean

$1/\mu_B$ hours. Suppose that arriving customers are willing to wait indefinitely if the worker is busy and if they are next in line. If an arriving customer sees another customer waiting, then she leaves the system. In other words, the queue has capacity 1 and customers do not leave without being served once they are in queue.⁴

For a pair of prices p_A and p_B for the two classes, since the arrival process is Poisson, the reward process till any time t is a renewal reward process. Each renewal cycle begins with the worker being idle, and ends when the worker finishes a job and there is no customer in queue. The long-run average rate of earning in this case can be computed to be,

$$\mathcal{R}(p_A, p_B) = \frac{\rho_A p_A (1 - p_A) + \rho_B p_B (1 - p_B)}{\frac{\lambda_A \mu_A (1 - p_A) + \lambda_B \mu_B (1 - p_B) + \mu_A \mu_B}{(\lambda_A (1 - p_A) + \lambda_B (1 - p_B) + \mu_A)(\lambda_A (1 - p_A) + \lambda_B (1 - p_B) + \mu_B)} + \rho_A (1 - p_A) + \rho_B (1 - p_B)}. \quad (13)$$

The details of this computation can be found in the Appendix. Suppose that we fix $\lambda_A = \mu_A = 1$ and set $\lambda_B = \mu_B = r$, so that $\rho_A = \rho_B = 1$. In Figure 1, I plot the optimal prices as r varies from 0 to 1 and then from 1 to 100.⁵ This allows us to see the effect on the optimal prices of changing the job length of class B while keeping the overall load the same as that of class A . Observe that when the expected job length of class B is larger than that of class A ($r < 1$), the optimal price p_B^* is larger than p_A^* . The inequality is reversed when $r > 1$, i.e., when class B jobs are shorter than class A jobs. This suggests that in these cases, the expected opportunity cost doesn't necessarily grow linearly in the length of the job. Intuitively, when the worker starts working on a job, the rejections of incoming jobs only begin when a new job arrives in the queue. Thus smaller jobs impose relatively smaller externalities per unit time since there is a greater chance that they finish before a new job arrives in the queue.




Figure 1 Optimal prices as a function of r .

⁴ If the queue has infinite capacity and customers wait indefinitely, then there are no externalities imposed by accepting any job and no price discrimination is necessary if the customers' per-unit-time willingness to pay is identically distributed.

⁵ The optimal prices were computed using the `scipy.optimize` package in Python.

3.2. A hybrid setting

On a related note to the setting described in the previous section, in platforms where there is a mixture of on-demand and patient customer classes, price discrimination across these classes is typically necessary even if the distribution of the customers' hourly willingness to pay and arrival/job-duration characteristics are identical. To see this, suppose that there are two classes of customers: On-demand (A) and Patient (B). On-demand customers have to be served immediately, and Patient customers are willing to wait as long as the average waiting time is finite (i.e., the queue doesn't explode). Moreover, preemption of a Patient job is allowed. Suppose that all job durations are exponentially distributed. The arrival rate and service rate of the On-demand class are $\lambda_A = \mu_A = 1$. The arrival rate and service rate of the Patient class are $\lambda_B = 1/\sqrt{2} - \epsilon$ and $\mu_B = 1$ for some small $\epsilon > 0$ (the choice may seem mysterious but will soon become clear). Suppose that the hourly valuations of the customers are distributed uniformly in $[0, 1]$.

Now clearly, as long as the effective residual service rate available for the Patient class due to time spent in serving the On-demand class is larger than λ_B , the Patient class imposes no externality on the system, since accepting a Patient customer incurs no loss in revenue to the worker. The optimal hourly price for the Patient class can then be computed to be 0.5. The question is whether the effective service rate for the Patient class is sufficient. Now, turning to the On-demand class, it is clear that accepting a job imposes an externality on the system. If this was the only class on the platform, then the optimal price is the one that maximizes,

$$\frac{x\bar{F}(x)}{1 + \bar{F}(x)} = \frac{x(1-x)}{2-x}.$$

The optimal price can be computed to be $p_A^* = 2 - \sqrt{2}$. At this price, the long-run fraction of time that the worker is not working on an On-demand job can be computed to be $1/(2 - p_A^*) = 1/\sqrt{2}$ (this is simply the ratio of the expected time until arrival of an On-demand job in a renewal cycle and the expected total length of a renewal cycle). Hence, the effective service rate available for the Patient customers is $1/\sqrt{2} \times 1$ which is more than λ_B . Thus we conclude that the optimal hourly price for the On-demand customers is $2 - \sqrt{2}$ and that for the Patient customers is 0.5.

This example suggests that in a general setup where there exist different classes of customers that are differentiated in their delay tolerance, which can be modeled by their job deadlines, price discrimination may be necessary for optimality despite identically distributed preferences.

3.3. Non-exponential discounting

One can show that price discrimination across customer classes because of differing arrival/duration characteristics may be necessary despite Poisson arrivals if the time discounting is non-exponential. We saw earlier that exponential discounting is equivalent to an exponentially distributed time horizon. Suppose instead that the time horizon is exponentially distributed with rate 1 with probability

1/2 and with rate 2 with probability 1/2. Suppose there are 2 classes of customers, A and B. The arrival rates are $\lambda_A = \lambda_B = 1$. The job durations are exponentially distributed with rates $\mu_A = 1$ and $\mu_B = 2$. Suppose that the hourly prices are p_A and p_B for the two classes. In the event that the time horizon is exponentially distributed with rate γ , the expected earning can be computed to be,

$$\frac{\hat{\rho}_A p_A (1 - p_A) + \hat{\rho}_B p_B (1 - p_B)}{1 + \hat{\rho}_A (1 - p_A) + \hat{\rho}_B (1 - p_B)}, \quad (14)$$

where $\hat{\rho}_A = \lambda_A / (\mu_A + \gamma)$ and $\hat{\rho}_B = \lambda_B / (\mu_B + \gamma)$. Thus the expected earnings can be computed to be,

$$\begin{aligned} \mathcal{R}(p_A, p_B) = & \frac{1}{2} \times \frac{(1/2)p_A(1 - p_A) + (1/3)p_B(1 - p_B)}{1 + (1/2)(1 - p_A) + (1/3)(1 - p_B)} \\ & + \frac{1}{2} \times \frac{(1/3)p_A(1 - p_A) + (1/4)p_B(1 - p_B)}{1 + (1/3)(1 - p_A) + (1/4)(1 - p_B)} \end{aligned} \quad (15)$$

The optimal prices can be computed to be $p_A^* \approx 0.56761$ and $p_B^* \approx 0.567089$.⁶ Thus in this case, price discrimination becomes necessary for optimality.

4. Price differentiation with heterogeneous customers

Coming back to the on-demand setting where there are no per-worker queues, it is now clear that under Poisson arrivals, 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 structural insights into the optimal prices in this case and finally present a procedure to compute these prices under certain standard assumptions. Let F_k be the distribution of hourly valuations of customers in class $k \in \mathcal{K}$, with support on $[0, \bar{v}_k]$. I assume that each F_k is differentiable on its support (with right and left derivatives defined at 0 and \bar{v}_k , respectively). Let $\bar{F}_k = 1 - F_k$ and f_k denote the corresponding tail distribution functions and densities, respectively. It is straightforward to establish as we did earlier (i.e., the derivation of (2)), that for a price vector \mathbf{p} , the long-run average earning is given by:

$$\mathcal{R}(\mathbf{p}) = \frac{\sum_k \rho_k (p_k - c) \bar{F}_k(p_k)}{1 + \sum_{k \in \mathcal{K}} \rho_k \bar{F}_k(p_k)}. \quad (16)$$

The problem of maximizing the long-run rate of earning is thus defined as:

$$\max_{\mathbf{p}: p_k \in [0, \bar{v}_k] \forall k} \mathcal{R}(\mathbf{p}). \quad (17)$$

⁶ Optimal prices computed using the `scipy.optimize` package in Python

Let the optimal rate of earning be denoted by \mathcal{R}^* . Then for any feasible set of prices \mathbf{p} , we have that

$$\begin{aligned}\mathcal{R}^* &\geq \frac{\sum_k \rho_k (p_k - c) \bar{F}_k(p_k)}{1 + \sum_{k \in \mathcal{K}} \rho_k \bar{F}_k(p_k)} \\ &\Rightarrow \mathcal{R}^* (1 + \sum_{k \in \mathcal{K}} \rho_k \bar{F}_k(p_k)) \geq \sum_k \rho_k (p_k - c) \bar{F}_k(p_k) \\ &\Rightarrow \mathcal{R}^* \geq \sum_k \rho_k (p_k - c - \mathcal{R}^*) \bar{F}_k(p_k).\end{aligned}\tag{18}$$

This expression is an equality if and only if \mathbf{p} is the vector of optimal prices \mathbf{p}^* (which need not be unique). This also implies that the optimal prices are the maximizers of the function $\sum_k \rho_k (p_k - c - \mathcal{R}^*) \bar{F}_k(p_k)$. Since this function is separable across the classes, this in turn implies that any optimal price for class k , p_k^* , maximizes $\bar{F}_k(p_k)(p_k - c - \mathcal{R}^*)$. This fact has an intuitive interpretation. The optimal earning rate \mathcal{R}^* is indeed the opportunity cost per hour of being busy. Thus, given \mathcal{R}^* , one simply solves the optimal pricing problem for each class assuming that the total hourly cost for serving that class is $c + \mathcal{R}^*$. This is a standard optimal pricing problem in mechanism design; c.f. Myerson (1981).

For a fixed hourly opportunity cost $\mathcal{R} \geq 0$, for each $k \in \mathcal{K}$, consider the optimization problem,

$$\max_{p \in [0, \bar{v}_k]} \bar{F}_k(p)(p - c - \mathcal{R}).\tag{19}$$

The first derivative of the objective function is given by,

$$\bar{F}_k(p) - (p - c - \mathcal{R})f_k(p) = -f_k(p) \left(p - c - \mathcal{R} - \frac{\bar{F}_k(p)}{f_k(p)} \right).\tag{20}$$

Now suppose that the function $p - \bar{F}_k(p)/f_k(p)$ is strictly increasing in $p \in [0, \bar{v}_k]$ for each $k \in \mathcal{K}$; if this holds we say that each distribution F is *strictly regular* (if $p - \bar{F}_k(p)/f_k(p)$ is non-decreasing then such a distribution is called *regular* in mechanism design literature; c.f. Myerson (1981)). For instance, this would hold if F_k has a non-decreasing hazard rate for each $k \in \mathcal{K}$, i.e., $f_k(p)/\bar{F}_k(p)$ is non-decreasing (this has been referred to as the *Monotone Hazard Rate* or *Increasing Failure Rate* assumption in the mechanism design literature).

If the distributions are strictly regular, the entire function in the parenthesis in equation (20) is increasing. Thus, from the first order optimality condition, we can conclude that the optimal value of p that solves (19) is unique: it is \bar{v}_k if $\bar{v}_k \leq c + \mathcal{R}$, otherwise it is the unique value of p that satisfies $p = c + \mathcal{R} + \bar{F}_k(p)/f_k(p)$. With some abuse of notation, let us denote this optimal value as $p_k(\mathcal{R})$. It is straightforward to verify that the function $p_k(\mathcal{R})$ defined on \mathbb{R}^+ is continuous, and moreover differentiable at all points except $\bar{v}_k - c$. Also, both left and right derivatives exist at $\bar{v}_k - c$. Further, $p'_k(\mathcal{R}) \geq 0$ for all $\mathcal{R} < \bar{v}_k - c$ (the optimal price is non-decreasing in \mathcal{R}) and

$p'_k(\mathcal{R}) = 0$ for all $\mathcal{R} > \bar{v}_k - c$. At $\bar{v}_k - c$, we have the left derivative $p'_{k-}(\bar{v}_k - c) \geq 0$ and the right derivative $p'_{k+}(\bar{v}_k - c) = 0$.

Thus, under the strict regularity assumption, we know that the optimal prices \mathbf{p}^* are unique and they are $(p_k(\mathcal{R}^*))_{k \in \mathcal{K}}$. The optimal earning rate \mathcal{R}^* is of course, unknown. But we know that \mathcal{R}^* satisfies the fixed point relation:

$$\mathcal{R}^* = \frac{\sum_{k \in \mathcal{K}} (p_k(\mathcal{R}^*) - c) \rho_k \bar{F}_k(p_k(\mathcal{R}^*))}{1 + \sum_{k \in \mathcal{K}} \rho_k \bar{F}_k(p_k(\mathcal{R}^*))}. \quad (21)$$

The following result shows that given the regularity assumption, 1) this fixed point is unique, and 2) it is the maximizer of the response function. These two facts are key to obtaining a natural procedure for iteratively computing \mathcal{R}^* .

LEMMA 2. *Suppose that the distributions F_k are strictly regular for each $k \in \mathcal{K}$. Then, there is a unique \mathcal{R}^* that satisfies the fixed point relation (21). Moreover, \mathcal{R}^* maximizes the function,*

$$M(\mathcal{R}) = \frac{\sum_{k \in \mathcal{K}} (p_k(\mathcal{R}) - c) \rho_k \bar{F}_k(p_k(\mathcal{R}))}{1 + \sum_{k \in \mathcal{K}} \rho_k \bar{F}_k(p_k(\mathcal{R}))}. \quad (22)$$

Proof. First, note that a fixed point always exists since the optimal earnings and the optimal set of prices define such a fixed point. We first establish that it is unique. It is straightforward to verify that the function $M(\mathcal{R})$ defined on \mathbb{R}^+ is continuous, and differentiable at all points except $\{\bar{v}_k - c; k \in \mathcal{K}\}$. Also at these points, both left and right derivatives exist (these properties follow from the corresponding properties of $p_k(\mathcal{R})$). Then at any point \mathcal{R} the right-derivative of $M(\mathcal{R})$ can be expressed as,

$$M'_+(\mathcal{R}) = \sum_{k \in \mathcal{K}} \frac{\partial M}{\partial p_k} p'_{k+}(\mathcal{R}) \quad (23)$$

$$= \sum_{k \in \mathcal{K}} \frac{-\rho_k f_k(p_k(\mathcal{R}))}{1 + \sum_{k' \in \mathcal{K}} \rho_{k'} \bar{F}_{k'}(p_{k'}(\mathcal{R}))} \left(p_k(\mathcal{R}) - c - M(\mathcal{R}) - \frac{\bar{F}_k(p_k(\mathcal{R}))}{f_k(p_k(\mathcal{R}))} \right) p'_{k+}(\mathcal{R}). \quad (24)$$

Here $p'_{k+}(\mathcal{R})$ is the right-derivative of p_k at \mathcal{R} . Let \mathcal{R}^* be a fixed point of M , i.e., $M(\mathcal{R}^*) = \mathcal{R}^*$. Then we have,

$$M'_+(\mathcal{R}^*) = \sum_{k \in \mathcal{K}} \frac{-\rho_k f_k(p_k(\mathcal{R}^*))}{1 + \sum_{k' \in \mathcal{K}} \rho_{k'} \bar{F}_{k'}(p_{k'}(\mathcal{R}^*))} \left(p_k(\mathcal{R}^*) - c - \mathcal{R}^* - \frac{\bar{F}_k(p_k(\mathcal{R}^*))}{f_k(p_k(\mathcal{R}^*))} \right) p'_{k+}(\mathcal{R}^*). \quad (25)$$

From the optimality conditions for $p_k(\mathcal{R}^*)$, we have that $p_k(\mathcal{R}^*) - c - \mathcal{R}^* - \bar{F}_k(p_k(\mathcal{R}^*)) / f_k(p_k(\mathcal{R}^*)) \leq 0$. Moreover, if $p_k(\mathcal{R}^*) - c - \mathcal{R}^* - \bar{F}_k(p_k(\mathcal{R}^*)) / f_k(p_k(\mathcal{R}^*)) < 0$ then $p'_{k+}(\mathcal{R}^*) = 0$. Thus $M'_+(\mathcal{R}^*) = 0$ for any \mathcal{R}^* such that $M(\mathcal{R}^*) = \mathcal{R}^*$.

This implies that for any \mathcal{R}^* , there is an ϵ such that for all \mathcal{R} that satisfy $\mathcal{R}^* < \mathcal{R} < \mathcal{R}^* + \epsilon$, $M(\mathcal{R}) < \mathcal{R}$. For such an \mathcal{R} , if $p_k(\mathcal{R}) - c - \mathcal{R} - \bar{F}_k(p_k(\mathcal{R})) / f_k(p_k(\mathcal{R})) = 0$, then $p_k(\mathcal{R}) - c - M(\mathcal{R}) -$

$\bar{F}_k(p_k(\mathcal{R}))f_k(p_k(\mathcal{R})) > 0$ and $p'_{k+}(\mathcal{R}) \geq 0$, while if $p_k(\mathcal{R}) - c - \mathcal{R} - \bar{F}_k(p_k(\mathcal{R}))/f_k(p_k(\mathcal{R})) < 0$, then $p'_{k+}(\mathcal{R}) = 0$. Thus, $M'_+(\mathcal{R}) \leq 0$. This implies that the function $M(\mathcal{R})$ is non-increasing after \mathcal{R}^* . Hence, there can be only one fixed point \mathcal{R}^* .

Next, we show that \mathcal{R}^* maximizes $M(\mathcal{R})$. This is straightforward, since for any \mathcal{R} , $M(\mathcal{R})$ is an achievable revenue rate, achieved by simply using the prices $(p_k(\mathcal{R}))_{k \in \mathcal{K}}$. Hence, $M(\mathcal{R}) \leq \mathcal{R}^* = M(\mathcal{R}^*)$. \square

We can now show the following main result of this section.

THEOREM 2. *Suppose that the distributions F_k are strictly regular for each $k \in \mathcal{K}$. Then starting from any $\mathcal{R}_0 \geq 0$, the sequence $(\mathcal{R}_t)_{t \in \mathbb{N}}$ obtained by the relation $\mathcal{R}_{t+1} = M(\mathcal{R}_t)$ converges to the unique fixed point \mathcal{R}^* of $M(\mathcal{R})$. \mathcal{R}^* is the optimal rate of earning and the corresponding prices $p_k(\mathcal{R}^*)$ for $k \in \mathcal{K}$ are the optimal prices.*

Proof. Since $M(\mathcal{R}) \leq \mathcal{R}^*$, for any \mathcal{R}_0 , we have that $\mathcal{R}_t \leq \mathcal{R}^*$ for $t \geq 1$. Moreover, $\mathcal{R}_{t+1} = M(\mathcal{R}_t) > \mathcal{R}_t$ for any $\mathcal{R}_t < \mathcal{R}^*$. Thus $(\mathcal{R}_t)_{t \in \mathbb{N}; t \geq 1}$ is a monotonically increasing sequence bounded by \mathcal{R}^* , and hence converges to some \mathcal{R}' by the monotone convergence theorem. It must be that $\mathcal{R}' = \mathcal{R}^*$, since if not, then $\mathcal{R}' < \mathcal{R}^*$ and hence $M(\mathcal{R}') > \mathcal{R}'$, which contradicts the fact that \mathcal{R}' is the limit point. \square

Example. Consider two customer classes with loads $\rho_1 = \rho_2 = 1$. The hourly willingness to pay of class 1 is uniformly distributed in $[0, 1]$, and that of class 2 is uniformly distributed in $[0, 2]$. Thus $\bar{F}_1(x) = (1 - x)\mathbb{1}_{\mathcal{R} \in [0, 1]}$ and $\bar{F}_2(x) = (1 - x/2)\mathbb{1}_{\mathcal{R} \in [0, 2]}$. Suppose that the hourly cost of service is 0. For a given hourly reserve price \mathcal{R} , $p_k(\mathcal{R})$ solves $\max(p - \mathcal{R})(1 - p)$ for $k = 1$ and $\max(p - \mathcal{R})(1 - p/2)$ for $k = 2$. We thus obtain $p_1(\mathcal{R}) = \frac{1+\mathcal{R}}{2}\mathbb{1}_{\mathcal{R} \in [0, 1]} + \mathbb{1}_{\mathcal{R} > 1}$ and $p_2(\mathcal{R}) = \frac{2+\mathcal{R}}{2}\mathbb{1}_{\mathcal{R} \in [0, 2]} + 2\mathbb{1}_{\mathcal{R} > 2}$.

Thus for $\mathcal{R} \leq 1$, we obtain,

$$M(\mathcal{R}) = \frac{(1 - \mathcal{R}^2)/4 + (4 - \mathcal{R}^2)/8}{1 + (2 - \mathcal{R})/4 + (1 - \mathcal{R})/2} = \frac{6 - 3\mathcal{R}^2}{2(8 - 3\mathcal{R})}.$$

And for $\mathcal{R} \in (1, 2]$, we obtain,

$$M(\mathcal{R}) = \frac{(4 - \mathcal{R}^2)/8}{1 + (2 - \mathcal{R})/4} = \frac{4 - \mathcal{R}^2}{2(6 - \mathcal{R})}.$$

$M(\mathcal{R}) = 0$ for $\mathcal{R} > 2$. The function $M(\mathcal{R})$ is plotted in Figure 2, where one can see that its unique fixed point is its maximizer. This is the point $\mathcal{R}^* \approx 0.40589$. Thus the optimal hourly prices are $p_1^* \approx (1 + 0.40589)/2 = 0.7029$ and $p_2^* \approx (2 + 0.40589)/2 = 1.2029$. Figure 3 shows the convergence of the iterative procedure to \mathcal{R}^* from different starting points. The convergence is quick (the plot shows 5 iterations). Observe that as expected, the output values are no larger than \mathcal{R}^* after the first iteration.




Figure 2 The plot of $y = M(\mathcal{R})$ and the line $y = \mathcal{R}$. The fixed point of $M(\mathcal{R})$ ($\mathcal{R}^* \approx 0.40589$) is also the maximizer of $M(\mathcal{R})$.




Figure 3 Convergence of the iterative procedure to $\mathcal{R}^* \approx 0.40589$ from different starting points

REMARK 1. The above procedure of iteratively estimating the earning rate and using it as the opportunity cost to compute the optimal prices for each customer class doesn't generally converge to the optimal prices in the presence of queueing, or in any system where price discrimination across classes is necessary despite identically distributed preferences. This can be seen from the example presented in Section 3.1. In that example, despite the customer classes having the identically distributed hourly valuations, I showed that price discrimination was necessary for optimality. But the above procedure can never converge to different prices for the two customer classes, since for any estimated earning rate, the optimal pricing problem solved by the procedure for the two classes in every iteration is identical.

5. Worker competition

A natural question that arises at this point is if any of these insights carry over in the presence of competition. Indeed, in reality, several workers offering the same services operate on any service platform. The traditional revenue management response to this concern would be that the customer arrival rates already factor in everything exogenous to the decisions of the worker in focus; in particular, these are the *residual* arrival rates after having accounted for the competition. However, such a response is valid only if one assumes that the competitive environment is stable, and that it does not change in response to the prices set by the worker. There is a priori no reason to assume so in our setting. In this section, I offer a rather optimistic resolution of this concern.

First, I start with a basic observation about a setting that I will refer to as *uniform workers setting*. Suppose that there are $N \geq 2$ workers that offer identical services, which are indistinguishable in terms of dimensions like quality, etc. Moreover, without loss of generality, suppose that their hourly service costs are identical and equal to 0. Also, suppose that there is a single customer class with Poisson arrivals with rate λ , and an arbitrary job length distribution with a finite mean

$1/\mu$. Each customer's maximum hourly willingness-to-pay is drawn independently from a strictly regular common distribution F . Moreover, suppose that each arriving customer, a) chooses the cheapest available worker whose hourly price is lower than her hourly willingness to pay, b) leaves if there is no such worker or if all the workers are unavailable, and c) breaks ties arbitrarily if there are multiple workers that offer the same smallest price. Given a vector of prices (p_1, \dots, p_N) chosen by the workers, let $\mathcal{R}_i(p_1, \dots, p_N)$ be the long-run average revenue of worker i for $i \in \{1, \dots, N\}$. Recall the definition of Nash equilibrium.

DEFINITION 1 (NASH EQUILIBRIUM). We say that the vector of prices (p_1, \dots, p_N) is a Nash equilibrium if for all $i \in \{1, \dots, N\}$,

$$\mathcal{R}_i(p_1, \dots, p_i, \dots, p_N) \geq \mathcal{R}_i(p_1, \dots, p'_i, \dots, p_N) \text{ for any } p'_i \in \mathbb{R}.$$

I informally call such a Nash equilibrium vector of prices a *price equilibrium*. The main observation is that in the uniform workers setting, there exists no price equilibrium. The basic idea is the same as the ‘‘Edgeworth Paradox’’ in economics Fisher (1925). We record this observation in the following proposition.

PROPOSITION 1. *In the uniform workers setting described above with $N \geq 2$, there exists no Nash equilibrium vector of prices.*

Proof. Suppose that there exists a price equilibrium and without loss of generality, let workers 1 and 2 offer the two smallest prices at equilibrium, $p_1 \leq p_2$. Let p^* be the unique (due to strict regularity of F) monopolistic optimal price when there is a single worker. First, note that it cannot be that $p_1 > p^*$, since if so then 1 can lower his price to p^* and make the monopolistic optimal revenue. Also, it cannot be that $p_1 < p^* \leq p_2$, since again 1 can raise her price closer to p^* and make a revenue closer to the monopolistic optimal revenue. Thus it must be that $p_1 \leq p_2 \leq p^*$. It cannot be that $p_1 < p_2 \leq p^*$, since 1 can raise her price while still being less than p_2 and obtain a higher revenue by inching closer to the monopolistic optimal price of p^* . Thus we have $p_1 = p_2 \leq p^*$; but again this cannot be an equilibrium unless $p_1 = p_2 = 0$, since either 1 or 2 can undercut the other and get a larger share of the market for negligible loss in revenue. But $p_1 = p_2 = 0$ cannot be an equilibrium either since either worker is incentivized to raise the price to capitalize (i.e., earn a positive expected revenue) in the situations where all the other workers are busy and they are the only worker available for the customer. Hence, there is no price equilibrium. \square

This suggests that amongst undifferentiated workers, there is no hope of obtaining a price equilibrium. However, it is rarely the case that the workers on a platform are undifferentiated – they are typically ranked in order of quality as determined by the reputation systems of the platforms. In this case, I argue that the situation may not be as bleak and a price equilibrium may exist.

Consider again the setting where there are K classes of customers, with Poisson arrivals (rate λ_k) and arbitrary job length distributions (mean $1/\mu_k$). Suppose that a class k customer's hourly willingness to pay is drawn independently from the distribution F_k . There are N workers, $\{1, \dots, N\}$, ranked in descending order of perceived quality, i.e., $1 > 2 > 3 > \dots > N$. Each worker i sets a price p_{ik} for a customer of class k . I propose the following model for how a customer accounts for price and quality information in making purchase decisions on platforms.

DEFINITION 2 (A BEST-I-CAN-AFFORD (BICA) BUYER). From the available workers, the customer chooses the highest ranked worker with an hourly price (for her class) less than or equal to her hourly willingness to pay. She leaves if there is no such worker.

I argue that this is a natural model for quality sensitive customer behavior on platforms. For example, for high-value software development jobs on labor platforms like Upwork, clients have predetermined management approved budgets for their hourly expenditures and naturally, they want to get the best available freelancer that meets their budget constraint. Moreover, in cases where customers are not expected to persist on the platform for long, the absolute values of the worker reputation scores may not be very informative beyond the ranking information that they provide and thus, opting for the best ranked worker that one can afford may be perceived to be a safe strategy in the face of quality uncertainty. Finally, the search and filtering abilities provided by platforms naturally support such behavior by making it easy to list products or services in decreasing order of quality scores.

With this model, I argue that not only does a price equilibrium exist, it can be computed efficiently. The idea is simple. Since the highest rated worker 1 gets priority from the customers, she is unaffected by the pricing decisions of the remaining workers. That is, when she is available, she gets priority from all the customers who can afford to pay the price that she sets, irrespective of prices set by other available workers. Thus, she can independently solve the single worker price optimization problem that we solved in Section 4. Once her prices are fixed, she consumes a part of the incoming demand of customers who can afford to pay her prices – not all since a fraction of these customers are rejected service from her because she is busy. The residual demand is thus composed of all the arriving customers who cannot afford to pay the first worker's prices and a fraction of the customers who are able to but are denied service due to unavailability. The second worker now has priority over these customers and is unaffected by the pricing decisions of the workers who have a lower rating than her. Thus, she again solves the single worker price optimization problem with the residual demand as earlier. Then the third worker sets the optimal prices with the residual demand that is rejected by the first and second workers and so on and so forth, until all the workers have set their prices. We record this observation in the following proposition.

PROPOSITION 2. *With workers ranked by their quality and BICA buyers, a price equilibrium exists.*

The above argument also implies the following two points. First, with customers' hourly willingness to pay identically distributed across multiple classes, price discrimination due to systemic reasons is unnecessary for worker 1, and assuming that worker 1 doesn't price discriminate, such a price discrimination is unnecessary for worker 2, and so on. Second, if the workers asynchronously implement the algorithm discussed in Section 4, this would lead to each worker converging to the equilibrium price. This is trivial to see for the highest rated worker. But once the highest rated worker's prices have converged, the second highest rated worker will face a stable environment and her prices will converge, and so on and so forth.

6. Discussion and future directions

Extensions. Below are a couple of extensions that are subsumed by the model.

1. **Hourly vs. per-job pricing.** Hourly pricing is equivalent to per-job pricing in terms of maximizing the long-run earnings assuming that a) the mean job duration is commonly known to the worker and the customer, and b) a customer's per-unit-time valuation is same as her per-job valuation divided by the mean duration. In this case, the optimal per job price is the same as the optimal hourly prices multiplied by the mean duration of the job.

2. **Commissions.** Typically, workers have to pay a fixed percentage of their revenue to the platform. In this case, if the worker charges p per hour to the customer, her net revenue is βp for some $\beta < 1$. In this case, we can simply redefine p_k to be the net revenue rate of the worker coming in from customer class k and the probability that an incoming customer is willing to support that rate can be defined to be $\bar{F}'_k(p) = \bar{F}_k(p/\beta)$. Thus the analysis remains identical.

Applicability. Removed from the context of on-demand platforms, the present paper essentially studies the optimal pricing problem faced by a single server in stochastic system with no queueing. The rising on-demand nature of services in the gig economy is the main reason why this question is of pertinence now and is likely the reason why this specific sub-case hasn't received much attention earlier.

As I mentioned in the introduction, I believe that the present formulation is the most applicable to on-demand labor markets for skilled or semi-skilled labor for jobs like graphic design, software development, application/website development etc. In these cases, the amount of time clients are willing to spend in hiring a worker is negligible compared to the duration of the job. Additionally, workers are not incentivized to multi-task, i.e., accept multiple jobs at the same time, for two reasons. One is that clients are time-sensitive and offering a short turnaround time is critical to winning a contract. Moreover, and perhaps more importantly, platforms like Upwork provide

worker monitoring service (by taking worker’s screen snapshots at random intervals) to the client to ensure that she is billed for exactly the hours that the workers spend working. This further diminishes the incentive to multi-task since there is no possibility of billing multiple clients for the same time spent. This also makes hourly billing mutually attractive to both the client and the worker as compared to fixed prices: the worker doesn’t have to worry about working too much without compensation (high uncertainty in duration is typical for large projects), and platform monitoring ensures that the client doesn’t overpay.

The present formulation can be applicable for pricing on rental marketplaces like Airbnb, Turo, etc., except for a couple of caveats. In these cases, there is seasonality in demand, e.g., weekends demand characteristics may be different from weekdays, summer demand may have different characteristics than other seasons, etc. Although these can be modeled as different customer classes, due to the cyclic nature of their arrivals, Poisson arrivals may not be an appropriate assumption. An additional concern is that hosts may want to price discriminate based on the time from the date of renting, similar to traditional revenue management settings, since high paying customers typically arrive later. Having said so, either because of the desire for simplicity of the price menu, or because these above two effects are believed to be negligible, the optimal pricing formulation presented in the paper can potentially be a reasonable approximation in these settings.

Future directions. There are several interesting avenues for further investigation. It would be interesting to quantify the impact of decentralized competitive pricing on platform revenues when compared to centralized pricing. Another issue is that of learning customer preferences. In the present work, I assumed that the worker knows the distribution of the customers’ per-unit-time willingness to pay. It would be interesting to explore how these distributions can be learned while strategically experimenting with prices. The novel aspect here is the need to account for externalities imposed by accepting jobs on the intertwined goals of revenue maximization and demand learning. I am optimistic that the structure of the optimal pricing solution and especially the iterative procedure that I presented in this paper will inform the design of good “learning while earning” pricing strategies in this setting.

References

- Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*, 37(3):668–691.
- Bakos, Y. and Katsamakas, E. (2008). Design and ownership of two-sided networks: Implications for internet platforms. *Journal of Management Information Systems*, 25(2):171–202.
- Banerjee, S., Johari, R., and Riquelme, C. (2015). Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC ’15.

- Bimpikis, K., Candogan, O., and Saban, D. (2019). Spatial pricing in ride-sharing networks. *Operations Research*.
- Birge, J., Candogan, O., Chen, H., and Saban, D. (2018). Optimal commissions and subscriptions in networked markets. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 613–614. ACM.
- Cachon, G. P., Daniels, K. M., and Lobel, R. (2017). The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, 19(3):368–384.
- Caillaud, B. and Jullien, B. (2003). Chicken & egg: Competition among intermediation service providers. *RAND journal of Economics*, pages 309–328.
- Caro, F. and Simchi-Levi, D. (2012). Optimal static pricing for a tree network. *Annals of Operations Research*, 196(1):137–152.
- Castillo, J. C., Knoepfle, D., and Weyl, G. (2017). Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 241–242. ACM.
- De Stefano, V. (2015). 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.
- Economides, N. and Katsamakas, E. (2006). Two-sided competition of proprietary vs. open source technology platforms and the implications for the software industry. *Management Science*, 52(7):1057–1071.
- Fisher, I. (1925). Edgeworth’s papers relating to political economy. *The Quarterly Journal of Economics*, 40(1):167–171.
- Gallager, R. G. (2013). *Stochastic processes: theory for applications*. Cambridge University Press.
- Hassin, R. and Haviv, M. (2003). *To queue or not to queue: Equilibrium behavior in queueing systems*, volume 59. Springer Science & Business Media.
- Lin, M., Li, S., and Whinston, A. B. (2011). Innovation and price competition in a two-sided market. *Journal of Management Information Systems*, 28(2):171–202.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.
- Olive, D. J. (2008). *Applied robust statistics*.
- Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the european economic association*, 1(4):990–1029.
- Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND journal of economics*, 37(3):645–667.
- Taylor, T. A. (2018). On-demand service platforms. *Manufacturing & Service Operations Management*, 20(4):704–720.
- Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review*, 100(4):1642–72.

Ziya, S., Ayhan, H., and Foley, R. D. (2006). Optimal prices for finite capacity queueing systems. *Operations Research Letters*, 34(2):214–218.

Ziya, S., Ayhan, H., and Foley, R. D. (2008). A note on optimal pricing for finite capacity queueing systems with multiple customer classes. *Naval Research Logistics (NRL)*, 55(5):412–418.

Appendix

A. Derivation of (13)

Let P_0 be the expected earnings between two successive times when the system is empty and the worker is idle (expected earnings in a renewal cycle). Let P_1 (P_2) be the expected earning till the system is empty starting from the time when a job of class A (B) is accepted and there is no customer in queue. Let T_0 be the expected time till the system is empty again after taking up a job, starting with an empty system (expected duration of a renewal cycle). Let T_1 (T_2) be the expected time till the system is empty starting from the time when a job of class A (B) is accepted and there is no customer in queue. Then from the renewal reward theorem, the long-run average earning converges almost surely to P_0/T_0 . Denote $\lambda'_i = \lambda_i(1 - p_i)$ for $i = 1, 2$. Then we have the following set of first-step equations satisfied by the various quantities.

$$P_0 = \frac{\lambda'_1 P_1}{\lambda'_1 + \lambda'_2} + \frac{\lambda'_2 P_2}{\lambda'_1 + \lambda'_2} \quad (26)$$

$$P_1 = \frac{p_1}{\mu_1} + \frac{\lambda'_1 P_1}{\lambda'_1 + \lambda'_2 + \mu_1} + \frac{\lambda'_2 P_2}{\lambda'_1 + \lambda'_2 + \mu_1} \quad (27)$$

$$P_2 = \frac{p_2}{\mu_2} + \frac{\lambda'_1 P_1}{\lambda'_1 + \lambda'_2 + \mu_2} + \frac{\lambda'_2 P_2}{\lambda'_1 + \lambda'_2 + \mu_2} \quad (28)$$

$$T_0 = \frac{1}{\lambda'_1 + \lambda'_2} + \frac{\lambda'_1 T_1}{\lambda'_1 + \lambda'_2} + \frac{\lambda'_2 T_2}{\lambda'_1 + \lambda'_2} \quad (29)$$

$$T_1 = \frac{1}{\mu_1} + \frac{\lambda'_1 T_1}{\lambda'_1 + \lambda'_2 + \mu_1} + \frac{\lambda'_2 T_2}{\lambda'_1 + \lambda'_2 + \mu_1} \quad (30)$$

$$T_2 = \frac{1}{\mu_2} + \frac{\lambda'_1 T_1}{\lambda'_1 + \lambda'_2 + \mu_2} + \frac{\lambda'_2 T_2}{\lambda'_1 + \lambda'_2 + \mu_2}. \quad (31)$$

Solving, we obtain (13) as the expression for P_0/T_0 .