

The Power of Two in Large Service-Marketplaces

Dheeraj Narasimha*

Srinivas Nomula[†]

Srinivas Shakkottai[‡]

Parimal Parag[†]

Abstract—We consider a large-scale service marketplace with numerous servers that scale with the job arrival rate. Jobs arrive with private valuations representing their willingness to pay. In a centralized system, jobs are matched to available servers, and prices are set in a centralized manner to maximize revenue. We investigate whether similar scalability can be achieved in a distributed marketplace where jobs are randomly matched to servers, which set their own prices based on job valuations and system occupancy. Our results show that matching a job to a single randomly selected server leads to increasing prices with the arrival rate, resulting in high blocking probability and reduced revenue. We then examine matching jobs to two servers, which compete to provide service if unoccupied. We demonstrate the existence and convergence to a mean field equilibrium (MFE) in this setup, where servers strategically respond to competitors' prices. We characterize the MFE and show that this two-server choice mitigates price scaling with the arrival rate, ensuring lower blocking probabilities and higher system revenue. Our findings are validated through simulations illustrating a variety of operating scenarios.

Index Terms—mean field games, parallel server systems, revenue maximization.

I. INTRODUCTION

The rise of service markets, especially in the area of cloud services, has significantly changed the way computational resources are allocated and utilized. Such servers support a variety of compute tasks, ranging from containers for microservices-based applications, machine learning, to blockchain mining and ledger handling. These markets are characterized by a large number of servers, owned by many different firms, which service a dynamic arrival of jobs, each with distinct valuations based on their urgency and computational needs. In such a high-demand environment, efficiently matching jobs to available servers while maximizing revenue and minimizing blocking probability is a significant challenge.

Traditional systems manage this problem by centralized matching of jobs to servers and centrally choosing a price for service. While this approach simplifies the pricing strategy, it

struggles to scale efficiently with increasing job arrival rates, since it needs to monitor each server and optimize prices across them all based on the demand in the form of job arrival rates, leading to higher complexity of market operations. As the demand for cloud services continually increases, there is a need for more scalable marketplace designs.

A distributed marketplace offers a promising solution. In this setup, jobs are randomly matched to servers, which then set their own prices based on the current system occupancy and the distribution of job valuations. This decentralized approach allows for more flexible pricing strategies that can better respond to fluctuating demand. However, it introduces its own set of complexities, particularly in ensuring that the system remains efficient and profitable as the scale increases.

With this context, we attempt to answer the question: *Can we design simple, randomized matching markets under which each server can competitively set its prices, while retaining high revenues and throughput of job completion?*

Main Results

We investigate the design of a large-scale service marketplace with a focus on two decisions, namely, (i) matching jobs to servers and (ii) pricing the service, and their impact on system performance. Under a centralized matching scheme, an incoming job can be matched to an unoccupied server, as long as one exists. This requires management by a central gateway that operates across all the servers, which is complex and less scalable. A simpler approach would be to simply match each incoming job with one or more randomly selected servers. This might result in a higher probability of jobs being blocked, thereby reducing the overall system revenue.

We consider pricing schemes under which the declared price for service is always a sample drawn from some exponential distribution. The reason for focusing on exponential distributions is that they provide a small amount of randomness over a fixed price, which prevents the system from getting stuck in corner cases, but the short-tailed nature of the exponential distribution implies that prices are close to the mean value. We study two pricing approaches. In the first, the parameter of the exponential distribution is chosen collaboratively across firms to maximize revenue in a baseline approach, whereas in the second, it is chosen competitively by servers belonging to different firms in a more realistic game setting.

We thus compare the following strategies across the decision rules for matching and pricing:

1. D_1C : Deterministic centralized matching of each job to one unoccupied server, and a collaborative choice of the price

Author* is with the Laboratoire d'Informatique de Grenoble, 38401, France. Authors[†] are with the Department of Electrical Communication Engineering, Indian Institute of Science, Bangalore, Karnataka 560012, India. Author[‡] is with the Department of Electrical and Computer Engineering and Department of Computer Science and Engineering, Texas A&M University, College Station, Texas 77843, USA.

This work is supported in part by National Science Foundation under Grant CNS-2312978; in part by Qualcomm Inc. under Qualcomm University Relations 6G India; in part by IBM Research Lab India under IBM-IISc Hybrid Cloud Lab (IHL) open research collaboration; in part by Anusandhan National Research Foundation (ANRF) under Grant CRG/2023/008854; in part by U.K.-India Education and Research Initiative (UKIERI) under Grant SPARC/2024/3927; and in part by Centre for Networked Intelligence (a Cisco Corporate Social Responsibility (CSR) Initiative), IISc. All opinions are those of the authors and do not necessarily represent those of the funding agencies.

parameter. This is the baseline approach that requires complete (unrealistic) collaboration across all firms.

2. R_2C : Random matching to two servers, but collaborative choice of the price parameter. This approach provides a baseline that eliminates the complexity of having a matching gateway that knows the occupancy of all the servers, but still requires collaboration in price-setting.

3. R_2G : Random matching to two servers, and competitive choice of price parameters under a strategic game setting. This approach eliminates the need for both centralized matching and collaborative pricing, and we study it as a mean field game (MFG) where servers have a belief about the distribution of prices of others and optimize their prices against that belief.

4. R_1C : Random matching to a single server and centralized choice of its price parameter. This approach eliminates the complexity of a centralized gateway, but there is no competition across servers. One would expect that this results in a higher probability of jobs being blocked, thereby reducing the overall system revenue.

We characterize and compare the revenue rate and throughput engendered in the above schemes. Our specific focus is on R_2G , which exploits the power of two random choices to increase revenue and throughput to values close to R_2C and attains a similar ultimate scaling as D_1C . Our main analytical contribution here is to demonstrate the existence and convergence to an MFE, wherein we have a consistent best-response price-setting policy and resultant price distribution.

Through our analysis, we show that this two-server matching strategy with competitive pricing mitigates the issue of price scaling with the arrival rate. As a result, it achieves lower blocking probabilities and higher system revenue compared to the single-server matching scenario of R_1C . Moreover, our approach provides a similar revenue scaling to the more complex strategies, but with a simpler implementation.

Our main methodological innovation is in showing convergence to the mean field limit in the R_2G setting. Existing approaches to showing MFE convergence rely on assuming that all agents follow the same policy. However, we develop a more robust approach wherein we explicitly model a “tagged server” that is an individual agent that plays strategically against an ensemble of other agents. We allow for tagged agent and the ensemble to each follow their own policies and are able to show strong exponential convergence of both to the MFE.

These findings are supported by numerical evaluations under different settings, highlighting the practical applicability of our proposed strategy. By demonstrating that a modest increase in the number of server choices can lead to substantial improvements in system performance, our study offers a pathway for more efficient and scalable service marketplaces.

Related Works

While dynamic pricing has been studied since the works of Cournot [1] further elaborated in [2], using pricing as a control mechanism can be traced back to [3]–[5]. For a more comprehensive view on dynamic pricing, we refer the reader to [6]–[8] to name a few.

More recently, there has been a focus on this form of pricing in the context of cloud computing [9]–[12] and the references therein as well as in the context of a multi-class multi server model in [13]. These models typically view the pricing scheme from the perspective of a central manager with limited servers.

Our work departs from the centralized perspective, instead looking at the problem as a game played by independent agents, see for example [14], [15] and their references. Such a view quickly becomes intractable since the computation of the revenue rate depends on the pricing policies of all the other agents. To circumvent this problem we shift to a mean field game problem. Here we view the game from the perspective of a typical agent against a continuum of agents that comprise the rest of the system. Such a view of games with many agents began with the works of [16] and [17] in parallel. Several works have looked to capitalize on this perspective for network applications [18]–[24] and networked markets [25], [26]. The former works deal with competition in resource constrained scenarios, while the latter works are closer to our setup on market design. However, they do not consider the service markets with competition across servers. Specifically, [26] considers one-to-one matching of consumers to producers in a prosumer marketplace, and so does not face the issues of blocking due to occupancy or benefit from the power of two choices in the manner of our system.

Like the above works, our goal is to exploit the analytical tractability of the mean field approach, while allowing us to numerically compute fixed points in terms of control policies for such complex systems. In doing so, we are able to obtain strong convergence results that are of independent interest, while addressing the specific problem of server matching and pricing in compute services markets.

II. SYSTEM MODEL

We consider a system with $N + 1$ servers labeled $\mathcal{N} \triangleq \{0, \dots, N\}$. Server 0 is called the tagged server, and remaining N servers in $[N] = \mathcal{N} \setminus \{0\}$ are untagged. Note, a mean field game approach consists of a typical agent viewing the rest of the system as the continuum limit of a homogeneous collection of agents with fixed strategy (known as the mean field limit) and proceeds to play a best response strategy. Each agent views themselves as a tagged server and the remaining servers as untagged, and hence by symmetry, the strategy is cloned across all servers. Tasks arrive into this system as a Poisson process with homogeneous rate $(N + 1)\lambda$. We denote the arrival instant of k th task by A_k , and the k th task inter-arrival time by $T_k \triangleq A_k - A_{k-1}$, for all $k \in \mathbb{N}$. Accordingly, the Poisson task arrival time sequence is denoted by $A \in \mathbb{R}_+^{\mathbb{N}}$, and the task inter-arrival time sequence by $T \in \mathbb{R}_+^{\mathbb{N}}$, where T is *i.i.d.* exponential with rate $(N + 1)\lambda$. We assume that the completion time for any task at any of the $N + 1$ servers is *i.i.d.* exponentially distributed random variable with rate 1. We assume that $S \in \mathbb{R}_+^{\mathbb{N}}$ is an *i.i.d.* exponentially distributed sequence with unit rate, where S_k is the service requirement of task k that arrives at instant A_k .

Each server has two states in $\mathcal{Z} \triangleq \{0, 1\}$, *busy* or *idle* denoted by 1 and 0 respectively. A server transitions from idle to busy when it accepts a task for service, and from busy to idle when it completes service. As such, we denote the state space of all servers by $\mathcal{X} \triangleq \mathcal{Z}^N$, where state of all servers $x \in \mathcal{X}$ is a vector and $x_n \in \mathcal{Z}$ indicates the occupancy of the server $n \in \mathcal{N}$. The state of the system is denoted by the process $X \in \mathcal{X}^{\mathbb{R}_+}$ where $X_t \in \mathcal{X}$ denotes the state of the system at time $t \in \mathbb{R}_+$. The random variable $X_{t,n} \in \mathcal{Z}$ indicates the occupancy of server $n \in \mathcal{N}$ at time $t \in \mathbb{R}_+$. We define the closed interval $\mathcal{Z} \triangleq [0, 1]$, and denote the fraction of the occupied untagged servers $\{1, \dots, N\}$ at time t by

$$Z_t^N \triangleq \frac{1}{N} \sum_{n=1}^N X_{t,n} \in \mathcal{Z}_N \triangleq \left\{0, \frac{1}{N}, \dots, 1\right\} \subseteq \mathcal{Z}.$$

A. Task valuation and server pricing

We denote the random valuation sequence for tasks by $V \in \mathbb{R}_+^N$, where each incoming task k has a random positive valuation $V_k \in \mathbb{R}_+$. We assume that V is *i.i.d.* with a common valuation distribution $G : \mathbb{R}_+ \rightarrow \mathcal{Z}$ and complimentary distribution denoted by $\bar{G} \triangleq 1 - G$. We assume that the valuation is exponentially distributed with rate v . Each server n sets a random and independent price $P_{k,n} : \Omega \rightarrow \mathbb{R}_+$ for each incoming task k . We assume that the tagged server 0 is a rational agent, and the rest of the servers are following the same random pricing function. The price distributions at the tagged server 0 and at each untagged server $n \in [N]$ are denoted by $F_0, F_1 : \mathbb{R}_+ \rightarrow \mathcal{Z}$ respectively. We consider the case when F_0, F_1 are exponentially distributed with rates d_0 and d_1 respectively.

B. Server selection and joining

Our approach to picking servers at random is inspired from a popular load-balancing method known as the power-of-two choices [27], [28]. Below, we adapt this method to our context. Each task k arrives at a dispatcher that picks a set of two servers $I_k \subseteq \mathcal{N}$ uniformly at random without replacement. If both servers in I_k are occupied, then the task k leaves the system having no effect. If one of the servers $n \in I_k$ is free, the task k probes the server for an asking price. If the task valuation $V_k > P_{k,n}$, it joins the idle server, if not, it once again leaves the system. Finally, if both servers in I_k are free, they must compete. The task k joins the lower priced of the two servers if its value lies above this price. That is, task k joins a server in I_k if $V_k > \min\{P_{k,n} : n \in I_k\}$. If both servers have an identical price, then it joins either of them with probability $\frac{1}{2}$ each. Otherwise, it joins the server offering a lower price. For $(N+1)$ server system, we denote the indicator of tagged server 0 being joined by task k as ξ_k^N . In terms of this indicator ξ_k^N , we can write the average revenue rate earned by the tagged server 0 during the first K task arrival times as

$$\bar{R}_{N,K} \triangleq \frac{1}{AK} \sum_{k=1}^K P_{k,0} \xi_k^N. \quad (1)$$

The performance metric of interest is the limiting revenue rate in the limit of large N , i.e.

$$\bar{R} \triangleq \lim_{N \rightarrow \infty} \lim_{K \rightarrow \infty} \bar{R}_{N,K}. \quad (2)$$

III. MEAN FIELD ANALYSIS

We look at the joint evolution $(X_{t,0}, Z_t^N, t \geq 0)$, where at time $t \in \mathbb{R}_+$, $X_{t,0}$ is the occupation indicator of the tagged server 0 and Z_t^N is the empirical distribution of occupied untagged servers in $[N]$. We denote the history of the process until time t as $\mathcal{F}_t^N \triangleq \sigma(X_{s,n}, n \in \mathcal{N}, s \leq t)$ for all $N+1$ servers. We denote the history of the process until the k th arrival time A_k as $\mathcal{F}_{A_k}^N$.

Definition 1. We define the indicator and probability of valuation of k th incoming task exceeding the price of the tagged server as

$$\eta_{k,10} \triangleq \mathbb{1}_{\{V_k > P_{k,0}\}}, \quad q_1 \triangleq \mathbb{E}\eta_{k,10}.$$

We define the indicator and probability of valuation of k th incoming task exceeding the price of the tagged server which is lower than another non-tagged server, as

$$\eta_{k,20} \triangleq \mathbb{1}_{\{V_k > P_{k,0}, P_{k,0} < P_{k,n}\}}, \quad q_{20} \triangleq \mathbb{E}\eta_{k,20}.$$

We define the indicator and probability of valuation of k th incoming task exceeding the price of a non-tagged server which is lower than the price of the tagged server, as

$$\eta_{k,2n} \triangleq \mathbb{1}_{\{V_k > P_{k,n}, P_{k,0} > P_{k,n}\}}, \quad q_{21} \triangleq \mathbb{E}\eta_{k,2n}.$$

We define the indicator and probability of valuation of k th incoming task exceeding the price of a single non-tagged server, as

$$\zeta_{k,1} \triangleq \mathbb{1}_{\{V_k > P_{k,n}\}}, \quad p_1 \triangleq \mathbb{E}\zeta_{k,1}.$$

We define the indicator and probability of valuation of k th incoming task exceeding the price of two non-tagged servers $n \neq m \in [N]$, as

$$\zeta_{k,2} \triangleq \mathbb{1}_{\{V_k > P_{k,n} \wedge P_{k,m}\}}, \quad p_2 \triangleq \mathbb{E}\zeta_{k,2}.$$

Lemma 1. For independent exponentially distributed pricing with rates d_0, d_1 for the tagged server 0 and remaining N servers respectively, and exponentially distributed valuation with rate v , we have

$$q_1 = \frac{d_0}{d_0 + v}, \quad q_{20} = \frac{d_0}{d_0 + d_1 + v}, \quad q_{21} = \frac{d_1}{d_0 + d_1 + v}, \\ p_1 = \frac{d_1}{d_1 + v}, \quad p_2 = \frac{2d_1}{2d_1 + v}.$$

Proof: The result follows from the fact that valuations and prices are independent exponential random variables. ■

Lemma 2. For the $(N+1)$ server system under consideration, the indicator of selection of tagged server 0 for service by the k th incoming task arrival is

$$\xi_k^N = \mathbb{1}_{\{0 \in I_k\}} \bar{X}_{A_k,0} \sum_{n=1}^N \mathbb{1}_{\{n \in I_k\}} \left(X_{A_k,n} \eta_{k,10} + \bar{X}_{A_k,n} \eta_{k,20} \right).$$

Proof: Tagged server 0 is selected for service by the k th incoming arrival, iff (a) it is selected in the random two-subset I_k for selection, (b) the tagged server is idle, (c) the value of the incoming arrival is higher than the price of the tagged server, and (c) either the other selected untagged server is busy or the selected untagged server is idle but has a higher price than the tagged server. ■

Recall that a randomly selected subset $I_k \in \binom{[N]}{2}$, and hence can have zero, one, or two occupied servers. Accordingly, we define the following probabilities.

Definition 2. We define the probability of one selected server being the tagged server and the occupancy of selected non-tagged server to be $j \in \mathbb{Z}, n \in [N]$, as

$$r_j \triangleq P\left(\cup_{n=1}^N \{I_k = \{0, n\}, X_{A_k, n} = j\} \mid \mathcal{F}_{A_k}^N\right).$$

We define the probability of none of the selected servers being the tagged server and the sum of the occupancy of two selected servers to be $\ell \in \{0, 1, 2\}$, as

$$s_\ell \triangleq P\left(\left\{0 \notin I_k, \sum_{n \in I_k} X_{A_k, n} = \ell\right\} \mid \mathcal{F}_{A_k}^N\right).$$

Remark 1. Summing over all possible states for two selected servers by the random selection I_k , we observe that $\sum_{j \in \mathbb{Z}} r_j = P\left(\{0 \in I_k\} \mid \mathcal{F}_{A_k}^N\right) = \frac{2}{N+1}$. Similarly, summing over the occupancy indicators of two selected servers, we obtain $\sum_{\ell \in \{0, 1, 2\}} s_\ell = P\left(\{0 \notin I_k\} \mid \mathcal{F}_{A_k}^N\right) = \frac{N-1}{N+1}$.

Lemma 3. Consider the selection probabilities defined in Definition 2 and let $Z_{A_k}^N = z$. The selection probabilities for the tagged server are

$$r_0 = \frac{2\bar{z}}{(N+1)}, \quad r_1 = \frac{2z}{(N+1)}.$$

The selection probabilities for non tagged servers are

$$s_0 = \bar{z} \frac{(N\bar{z} - 1)}{N+1}, \quad s_1 = \frac{2N\bar{z}z}{(N+1)}, \quad s_2 = z \frac{(Nz - 1)}{N+1}.$$

Proof: Recall that $\sigma(X_{A_k, n}) \subseteq \mathcal{F}_{A_k}^N$ for all $n \in \mathcal{N}$ and I_k is selected independently and uniformly at random at each arrival instant. In particular, $P\left(\{I_k = \{0, n\}\} \mid \mathcal{F}_{A_k}^N\right) = \frac{2}{N(N+1)}$ for all $n \in [N]$, and hence from the the law of total probability, we can write

$$r_j = \frac{2}{N(N+1)} \sum_{n=1}^N \mathbb{1}_{\{X_{A_k, n} = j\}} = \frac{2}{N(N+1)} (jNz + \bar{j}N\bar{z}).$$

Similarly, using the law of total probability for disjoint events and the fact that I_k is chosen uniformly at random, we can write

$$s_\ell = \frac{2}{N(N+1)} \sum_{n=1}^N \sum_{m=1}^{n-1} \mathbb{1}_{\{X_{A_k, n} + X_{A_k, m} = \ell\}}.$$

We observe that $\mathbb{1}_{\{X_{A_k, n} + X_{A_k, m} = 0\}} = \bar{X}_{A_k, n} \bar{X}_{A_k, m}$ and $\mathbb{1}_{\{X_{A_k, n} + X_{A_k, m} = 2\}} = X_{A_k, n} X_{A_k, m}$. Results follow from the definition of $Z_{A_k}^N$. ■

A. System evolution

Proposition 1. The process $(X_{t,0}, Z_t^N, t \geq 0) : \Omega \rightarrow (\mathbb{Z} \times \mathbb{Z}_N)^{\mathbb{R}_+}$ is a continuous time Markov chain with the generator matrix Q^N defined for all $x, y \in \mathbb{Z}$ and $z, w \in \mathbb{Z}_N$ as

$$\begin{aligned} Q_{(x,z),(y,w)}^N &= \left(Nz \mathbb{1}_{\{w=z-\frac{1}{N}\}} + \lambda \bar{z} [2p_1(x + Nz) \right. \\ &\quad \left. + 2\bar{x}q_{21} + p_2(N\bar{z} - 1)] \mathbb{1}_{\{w=z+\frac{1}{N}\}} \right) \mathbb{1}_{\{y=x\}} \\ &\quad + (x \mathbb{1}_{\{y=x-1\}} + 2\lambda \bar{x}(zq_{11} + \bar{z}q_{20}) \mathbb{1}_{\{y=x+1\}}) \mathbb{1}_{\{w=z\}}. \end{aligned}$$

This CTMC is irreducible and has finite states and hence positive recurrent with stationary distribution $\pi^N \in \mathcal{M}(\mathbb{Z} \times \mathbb{Z}_N)$.

Proof: The result can be derived from the fact that arrivals are Poisson, service times are i.i.d. exponentially distributed, and the prices and valuations are independent exponential random variables. Proof is omitted for brevity. ■

B. McKean-Vlasov equation

If the empirical distribution of states of an N -particle stochastic system is Markov, then it can be shown that under certain conditions this stochastic evolution converges to a deterministic evolution for large N , which is the limit of the empirical dynamic over N particles. This deterministic evolution is characterized by McKean-Vlasov equation. For the proposed system, this equation is defined in Definition 4.

Definition 3. Consider a family of CTMCs with generator matrix Q^N for each $N \in \mathbb{N}$. For each Q^N , we define an associated map $h_N : \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Z}$ defined for any $(x, z) \in \mathbb{Z} \times \mathbb{Z}$, as

$$h_N(x, z) \triangleq \sum_{(y,w) \in \mathbb{Z} \times \mathbb{Z}_N} Q_{(x,z),(y,w)}^N (w - z).$$

For this family of CTMCs we define the point-wise limit $h \triangleq \lim_{N \rightarrow \infty} h_N$.

Lemma 4. For the family of CTMCs with generator matrix Q^N in Proposition 1, the associated limiting map $h : \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Z}$ defined in Definition 3, is given by

$$h(x, z) = \lambda \bar{z} (2zp_1 + \bar{z}p_2) - z, \quad (x, z) \in \mathbb{Z} \times \mathbb{Z}. \quad (3)$$

Proof: From the definition of h_N and generator matrix Q^N in Definition 3 and Proposition 1 respectively, we obtain

$$h_N(x, z) = \frac{\lambda \bar{z}}{N} \left[(2p_1 - p_2)x + (2q_{21} - p_2)\bar{x} \right] + h(x, z).$$

We obtain the result by taking the point-wise limit of h_N as $N \rightarrow \infty$. ■

Remark 2. We observe that the $h(x, z)$ doesn't depend on x , and hence we will denote this map by $h : \mathbb{Z} \rightarrow \mathbb{Z}$. The absolute difference between maps h and h_N for each $(x, z) \in \mathbb{Z} \times \mathbb{Z}$ is

$$|h(z) - h_N(x, z)| = \frac{\lambda \bar{z}}{N} |x(2p_1 - p_2) + \bar{x}(2q_{21} - p_2)|. \quad (4)$$

Definition 4. [McKean-Vlasov equation] In terms of the map $h : \mathbb{Z} \rightarrow \mathbb{R}_+$ defined in (3), we can define an autonomous

nonlinear system $\Phi : \mathbb{R}_+^2 \times \mathcal{Z} \rightarrow \mathcal{Z}$ for $t \geq \tau \in \mathbb{R}_+$ and $z \in \mathcal{Z}$ as

$$\Phi_{t,\tau}(z) \triangleq z + \int_{\tau}^t h(\Phi_{s,\tau}(z)) ds, \quad \Phi_{\tau,\tau}(z) = z. \quad (5)$$

This equation is referred to as McKean-Vlasov equation and determines the evolution of the deterministic mean field model. For autonomous nonlinear system Φ defined in Definition 4, we define the set of rest points as $z^* \triangleq \{z \in \mathcal{Z} : h(z) = 0\}$.

Lemma 5. *For exponential distribution for valuation with rate v and pricing for untagged servers with rate d_1 , there is a unique rest point z^* for the deterministic mean field model, written in terms of $\alpha \triangleq \frac{v}{d_1}$, as*

$$z^* \triangleq -\frac{1}{2} \left(\alpha + \frac{(1+\alpha)(2+\alpha)}{2\lambda} \right) + \sqrt{1 + \alpha + \frac{1}{4} \left(\alpha + \frac{(1+\alpha)(2+\alpha)}{2\lambda} \right)^2}. \quad (6)$$

Proof: Any rest point $z \in \mathcal{S}$ of the deterministic mean field equation defined in Definition 4 satisfies $h(z) = 0$. From Lemma 4, we know the form of h , and rearranging the equation, we obtain that a rest point is the solution of a quadratic equation. Since $z \in \mathcal{Z}$, it follows that z^* is the positive root of this quadratic equation. The result follows from definition of p_1, p_2, α . Details are omitted for brevity. ■

C. Mean field convergence

We are now ready to show our main result that the equilibrium of the empirical Markov process converges to the rest point (6) of the deterministic evolution given by McKean-Vlasov equations, together with the associated convergence rate. To this end, we define a centred dynamical system and some constants which will improve the presentation of the convergence theorem, followed by the convergence result. A visual representation of this convergence for a specific numerical example is plotted in Fig. 1.

Definition 5. We define the distance of the system from a rest point z^* as $\varepsilon : \mathbb{R}_+ \times \mathcal{Z} \rightarrow \mathbb{R}_+$ for initial time $\tau = 0$, from any initial point $z \in \mathcal{Z}$, and at any time $t \in \mathbb{R}_+$, as

$$\varepsilon(t, z) \triangleq \Phi_{t,0}(z) - z^*. \quad (7)$$

Definition 6. We define the following two positive constants

$$L_\lambda \triangleq 2\lambda(2p_1 - p_2), \quad K_\lambda \triangleq 1 + 2\lambda(p_2 - p_1).$$

Theorem 1. *Let $z^* \in \mathcal{Z}$ be the unique rest point of the mean field model (5). The stationary fraction of occupied tagged servers Z_∞^N converges in the mean-square sense with rate $1/N$, i.e. $\mathbb{E} |Z_\infty^N - z^*|^2 = O\left(\frac{1}{N}\right)$.*

Remark 3. We will show that Theorem 1 holds under the following four conditions and then verify that these conditions are satisfied under our problem setup.

1) **Bounded mean transition-rate condition.** Stationary mean rate for transitions for generator matrix Q^N defined in Proposition 1 are bounded, i.e.

$$\mathbb{E} \sum_{(y,w) \in \mathcal{Z} \times \mathcal{Z}_N} Q_{(X_{\infty,0}, Z_\infty^N), (y,w)}^N |w - Z_\infty^N|^2 \leq \frac{1}{N}. \quad (8)$$

2) **Asymptotically accurate mean field model.** The mean field model defined in Definition 4 is asymptotically accurate mean field model, i.e.

$$\sup_{(x,z) \in \mathcal{Z} \times \mathcal{Z}_N} |h(x, z) - h_N(x, z)| \leq \frac{\lambda}{N}. \quad (9)$$

3) **Lipschitz partial derivatives.** The derivative $h'(z)$ exists and is L_λ Lipschitz for all $z \in \mathcal{Z}$. Further

$$K_\lambda \leq |h'(z)| \leq K_\lambda + L_\lambda. \quad (10)$$

4) **Global exponential stability.** The dynamical system $\dot{z} = h(z)$ for h defined in Lemma 4 is globally exponentially stable. In particular, for any initial condition $z_0 \in \mathcal{Z}$ and time $t \in \mathbb{R}_+$, the distance between solution of the McKean-Vlasov equation and its rest point is exponentially upper bounded as

$$|\varepsilon(t, z_0)| \leq |\varepsilon(0, z_0)| e^{-\lambda p_2 t}. \quad (11)$$

Typically such theorems have condition 1 split into separate conditions on the boundedness of the transition rates and the boundedness of the difference of states. However, since we keep the additional state of the tagged server, we needed this combined condition. Our proof gets simplified due to the global exponential stability of the rest point, which is typically not true in many mean field convergence cases. See [29] for an elaboration on the conditions.




Fig. 1: Convergence to the fixed point z^* as N grows larger for $\lambda = 0.5, v = 0.1, d_0 = 0.2$ and $d_1 = 0.9$ from $[10^5, 10^5 + 500]$ iterations. From, (6), $z^* = 0.975$. As N increases, the mean occupancy Z_t^N moves towards the rest point.

D. Equilibrium at the tagged server

We have studied the joint evolution of occupancy $X_{t,0}$ at the tagged server 0 and the fraction Z_t^N of occupied untagged servers $[N]$. We have shown that the stationary fraction of occupied untagged servers converges in mean sense to z^*

for large N . In the following theorem, we show that the joint distribution of $(X_{\infty,0}, Z_{\infty}^N)$ converges to an explicit distribution for large N . This follows from the fact that at time stationarity for large N , the occupancy evolution of the tagged server is a two-state continuous time Markov chain with transition rates dependent on the deterministic limit z^* .

Theorem 2. *Let z^* be the unique rest point of the mean field model (5). The stationary distribution π^N of the CTMC $(X_{t,0}, Z_t^N : t \geq 0)$ converges point-wise for all (x, z) , i.e.,*

$$\lim_{N \rightarrow \infty} \pi_{x,z}^N = \pi_{x,z} = \mathbb{1}_{\{z=z^*\}} \frac{\bar{x} + 2x\lambda(zq_1 + \bar{z}q_{20})}{2\lambda(zq_1 + \bar{z}q_{20}) + 1}. \quad (12)$$

Proof: From Theorem 1, the equilibrium fraction of occupied tagged servers Z_{∞}^N converges to z^* in the mean square sense, and hence also in distribution. In particular, we have $\lim_{N \rightarrow \infty} \sum_{x \in \mathcal{Z}} \pi_{x,z}^N = \mathbb{1}_{\{z=z^*\}} = \sum_{x \in \mathcal{Z}} \pi_{x,z}$. For each current state $(y, w) \in \mathcal{Z} \times \mathcal{Z}_N$, the possible transitions are from possible previous states $(x, z) \in \{(y, w), (y, w - \frac{1}{N}), (y, w + \frac{1}{N}), (\bar{y}, w)\}$ where $(x, z) = (y, w)$ is the self transition. Further, we can write $\pi_{x,z} = \mathbb{1}_{\{z=z^*\}} \pi_{x|z^*}$, to obtain that $\sum_{(x,z)} \pi_{x,z} Q_{(x,z),(y,w)}^N$ equals

$$\mathbb{1}_{\{z=z^*\}} \left(\pi_{\bar{y}|z^*} Q_{(\bar{y},z^*), (y,z^*)}^N - \pi_{y|z^*} Q_{(y,z^*), (\bar{y},z^*)}^N \right).$$

From the form of Q^N in Proposition 1 and the given form for $\pi_{x|z^*}$ in (12), we obtain that $\pi Q^N = 0$ for all $N \in \mathbb{N}$. ■

IV. EXISTENCE OF MEAN FIELD GAME EQUILIBRIUM

In Section III, we established that if we consider a system with N identical untagged servers with price parameter d_1 and a tagged server with price parameter d_0 , the fraction of busy servers converges to a point-mass $z^*(d_1)$ at rate $O(\frac{1}{N})$. Note, here we are assuming that the value of the incoming tasks are *i.i.d.* exponential with rate v . In this section, we examine a best response dynamic taken by the tagged server at the deterministic occupancy $z^*(d_1)$ of N untagged servers.

Given the tagged server's best response rate $d_0(z^*, d_1)$, the N untagged servers will adopt this response $d'_1 = d_0(z^*, d_1)$ to establish a new mean field $z^*(d'_1)$, in turn entailing a new best response from our tagged server. This naturally prompts the following question. Does a fixed point to such a dynamic exist? We identify such fixed points as the mean field game equilibrium defined below.

Definition 7. For a fixed value rate v , the mean field game equilibrium is defined by the pairwise fixed point in empirical measure and price parameter, (z_M, d_M) given by

$$z_M = z^*(d_M), \quad d_M = d_0(z_M, d_M). \quad (13)$$

At the end of this section, we will show the existence of the mean field game equilibrium under certain conditions.

A. Best response

In order to analyze the best response of the tagged server, we begin by quantifying average revenue rate (2) earned by this server.

Theorem 3. *In the large server limit, the limiting time average of the revenue rate at the tagged server 0 is*

$$\bar{R} = \lim_{N \rightarrow \infty} \lim_{K \rightarrow \infty} \bar{R}_{N,K} = \frac{z^* q_1^2 + \bar{z}^* q_{20}^2}{d_0(\frac{1}{2\lambda} + z^* q_1 + \bar{z}^* q_{20})}. \quad (14)$$

Proof: Consider the positive recurrent CTMC $(X_{t,0}, Z_t^N, t \geq 0)$ with generator matrix Q^N defined in Proposition 1, and associated invariant distribution π^N . We can rewrite the time average of the revenue rate as $\bar{R}_{N,K} = \frac{K}{A_K} \left(\frac{1}{K} \sum_{k=1}^K P_{k,0} \xi_k^N \right)$. From the strong law of large numbers, it follows that $\lim_{K \rightarrow \infty} \frac{K}{A_K} = (N+1)\lambda$ almost surely for an $N+1$ server system. Recall that I_k is selected uniformly among all 2-sets of $\mathcal{N}$, and the valuation and pricing are independent exponential random variables with rates v and d_1 , respectively. Therefore, from the explicit form in Lemma 2 for the selection indicator ξ_k^N of tagged server for service by the k th incoming task, we can write its conditional mean $\mathbb{E}[\xi_k^N | X_{A_k,0}, Z_{A_k}^N, P_{k,0}]$ given the states $X_{A_k,0}, Z_{A_k}^N$ and price $P_{k,0}$ at the tagged server, as

$$\frac{2\bar{X}_{A_k,0}}{N+1} \left(Z_{A_k}^N e^{-vP_{k,0}} + \bar{Z}_{A_k}^N e^{-vP_{k,0}} e^{-d_1 P_{k,0}} \right).$$

From the tower property of conditional expectation, the fact that $P_{k,0}$ is an exponential random variable with rate d_0 , and the definition of q_1, q_{20} from Lemma 1, we obtain

$$\mathbb{E}[P_{k,0} \xi_k^N | X_{A_k,0}, Z_{A_k}^N] = \frac{2\bar{X}_{A_k,0}}{d_0(N+1)} \left(Z_{A_k}^N q_1^2 + \bar{Z}_{A_k}^N q_{20}^2 \right).$$

From the PASTA [30] property, we know that Poisson arrival sees time averages. Therefore,

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \bar{X}_{A_k,0} Z_{A_k}^N = \sum_{z \in \mathcal{Z}_N} z \pi_{0,z}^N.$$

The result follows from taking limit $N \rightarrow \infty$ and using the expression for invariant distribution π from Theorem 2. ■




Fig. 2: Convergence of revenue rate to the rest point for $v = 0.1, d_0 = 0.2, d_1 = 0.9$ over 10^6 iterations.

Remark 4. Substituting for q_1, q_{20} in the revenue rate expression, we can verify that if $d_0 = 0$ or $d_0 = \infty$, then the average revenue rate is 0. Most importantly, we note that the deterministic fraction z^* of occupied servers only depends on price parameter d_1 of untagged servers, normalized arrival rate

λ , and the value parameter v . Hence, if we fix the policy of all untagged servers, then the revenue rate for the tagged server becomes a *deterministic* function of its price parameter d_0 . Thus, we can characterize the best response of the tagged server by finding the price parameter that maximizes its revenue rate. We plot the time convergence of the revenue rate for different numbers of servers in Fig. 2.

B. Unimodality

For fixed normalized arrival rate λ , value parameter v , and price parameter d_1 for untagged servers, let $\bar{R}(d_0)$ denote the revenue rate for the tagged server as a function of its price parameter d_0 . The following theorem uses the closed-form expression derived in (14) to characterize the shape of $\bar{R}(d_0)$.

Theorem 4. *For fixed d_1, v, λ and under the large server limit, the limiting revenue rate for the tagged server is unimodal for its price parameter $d_0 \in (0, \frac{v}{\sqrt{3}-1}]$, and hence there is a unique best response*

$$d_0^* \triangleq \arg \max \left\{ \bar{R}(d_0) : d_0 \in \left(0, \frac{v}{\sqrt{3}-1}\right] \right\}. \quad (15)$$

Further, $\bar{R}(d_0)$ is monotone decreasing for $d_0 \geq 2v$.

Remark 5. We note that restricting optimal price parameter $d_0^* \leq \frac{v}{\sqrt{3}-1}$ might seem unrealistic. However, the second part of our theorem shows that a maximizer of $\bar{R}(d_0)$ can only lie in $(0, 2v]$ which partially justifies this restriction. Further, we conducted numerical experiments for a range of system parameters λ, v, d_1 , and it appears that the limiting revenue rate is *unimodal* over the entire domain. We have plotted the limiting revenue rate for an example set of system parameters in Fig. 3. However, we note that we currently do not have analytical unimodality results for $d_0 \in (\frac{v}{\sqrt{3}-1}, 2v]$. One of the consequences of Theorem 4 is that there is a unique best response price parameter d_0^* for the tagged server.




Fig. 3: Limiting revenue rate $\bar{R}$ for the tagged server as a function of its price parameter $d_0 \in [0, 2v]$ for fixed system parameters $v = 0.5, \lambda = 1$ and different values of d_1 .

C. Existence of mean field game equilibrium

Having shown the unimodality of the limiting revenue rate, we next show the existence of mean field game equilibrium for our system.

Theorem 5. *For the $(N + 1)$ server system under consideration, we let the tagged server 0 select its best response price*

parameter d_0^ defined in (15) for a price parameter d_1 for the untagged servers. In the large server limit, there exists at least one mean field game equilibrium (z_M, d_M) as defined in Definition 7.*

Proof: We fix the value parameter v and normalized arrival rate λ . Given the price parameter d_1 chosen by all N untagged servers, we know that there exists a unique fraction of busy servers $z^*(d_1)$ by Theorem 1 at stationarity under large N limit. Further, this map $d_1 \mapsto z^*$ is continuous from Lemma 4. Next, it follows from Theorem 4 that there is a unique best response price parameter $d_0^* \in [0, \frac{v}{\sqrt{3}-1}]$ for the tagged server. Further, it can be verified that $\bar{R}(d_0)$ is continuous in d_0 and differentiable for $d_0 > 0$. From Berge's Maximum Theorem [31] it follows that d_0^* is a continuous function of z^* . Therefore, the map $z^* \mapsto d_0^*$ is continuous. This composition of maps can be summarized by the relation $d_0 \mapsto z^* \mapsto d_0^*$. Since this is a composition of two continuous maps, the map $d_0 \mapsto d_0^*$ is also continuous. It follows from Brower's fixed point theorem [32], that there exists at least one fixed point for this composition of maps and we obtain the desired result. ■

Remark 6. The mean field limit z^* is computable in closed-form and given in Lemma 4. Gradient ascent allows us to numerically compute the best response price parameter due to the unimodality of the limiting revenue rate. However, it is not straightforward to find an analytical closed-form expression for the best response. Hence, we can only numerically evaluate the mean field game equilibrium given value parameter v and normalized arrival rate λ .

V. COMPARISON WITH OTHER PRICING POLICIES

We measure performance in terms of the following three per-server metrics (a) limiting revenue rate $\bar{R}$, (b) mean price $\bar{P}$, and (c) throughput. If p_b is the probability of blocking an incoming arrival, then the throughput, ρ , is given by $\lambda(1 - p_b)$ for our systems. In this section, we compare the performance of our system R_2G to other systems D_1C, R_2C, R_1C mentioned in the introduction, in terms of these three metrics. We first compute the limiting revenue rate of the tagged server given the price parameter of untagged servers.

Lemma 6. *For a fixed v, λ and $d_0 = d_1 = d$, we can write the limiting occupancy of $(N + 1)$ server system as $z^*(d)$, probabilities $q_1(d), q_{20}(d)$. For this system, the limiting revenue rate is*

$$\bar{R}(d) \triangleq \frac{z(d)q_1^2(d) + (1 - z(d))q_{20}^2(d)}{d(\frac{1}{2\lambda} + z(d)q_1(d) + (1 - z(d))q_{20}(d))}, \quad (16)$$

and the blocking probability of an incoming arrival is

$$p_b(d) \triangleq \left(z^2(d) + \frac{2z(d)\bar{z}(d)v}{v + d} + \frac{\bar{z}^2(d)v}{v + 2d} \right). \quad (17)$$

Remark 7. We do not know whether the map $d \mapsto \bar{R}$ is unimodal, and hence it is not clear whether there is a unique maximizer for the limiting revenue rate. Thus, we only study this function numerically, as shown in Figure 4.




Fig. 4: Unimodality of limiting revenue rate as a function of common price parameter d for normalized arrival rate $\lambda = 1$.

At mean field game equilibrium (z_M, d_M) of Definition 7, all servers have the identical price parameter d_M and $\pi_{1,z} = z_M \mathbb{1}_{\{z=z_M\}}$, and hence have identical limiting revenue rates.

Corollary 1 (R₂G). Consider System R₂G at the mean field game equilibrium (z_M, d_M) in Definition 7. In terms of this pair and the limiting revenue rate $\bar{R}(d)$ from Lemma 6, the limiting revenue rate, the mean price, and throughput at each server are

$$\bar{R}_{R_2G} \triangleq \bar{R}(d_M), \quad \bar{P}_{R_2G} \triangleq \frac{1}{d_M}, \quad \rho_{R_2G} \triangleq \lambda \bar{p}_b(d_M). \quad (18)$$

System R₂C differs from System R₂G in that the identical price parameter at all servers is chosen to maximize the revenue rate $\bar{R}_{R_2C}$.

Corollary 2 (R₂C). Consider System R₂C with the following unique maximizer

$$d^* \triangleq \arg \max \{ \bar{R}(d) : d \in \mathbb{R}_+ \}. \quad (19)$$

Then the limiting revenue rate, mean price, and throughput at each server are

$$\bar{R}_{R_2C} \triangleq \bar{R}(d^*), \quad \bar{P}_{R_2C} \triangleq \frac{1}{d^*}, \quad \rho_{R_2C} \triangleq \lambda \bar{p}_b(d^*). \quad (20)$$

Remark 8. Consider a variant of System D₁C with constant uniform price p for all incoming arrivals. It follows that thinned Poisson arrival rate to the system is $(N+1)\lambda\bar{G}(p)$. For normalized arrival rate λ and price p such that $\lambda\bar{G}(p) < 1$. Then, it follows from [10], [33] that under a large server limit with centralized routing to idle servers, the following uniform price maximizes the limiting revenue rate

$$p_U \triangleq \arg \max_p p\bar{G}(p). \quad (21)$$

For exponential valuation, the maximizing uniform price $p_U = \frac{1}{v}$ for $\lambda < e$. Let us consider $\lambda \geq e$ and uniform price p such that $\lambda\bar{G}(p) < 1$, then the limiting revenue rate maximizing uniform price is

$$p_U \triangleq \arg \max_p \{ pe^{-vp} : e \leq \lambda < e^{vp} \} = \frac{1}{v} \ln \lambda.$$

Remark 9. We note that the only difference between the limiting revenue rate maximizing system and System D₁C

are the constant price in first system and *i.i.d.* exponentially random pricing in the second system. We will show that the qualitative behavior of System D₁C is very similar to the limiting revenue rate maximizing system in terms of the performance metrics under consideration. We have adapted the limiting revenue rate maximizing system to System D₁C to be consistent with random exponential pricing considered in all comparison systems.

Proposition 2 (D₁C). For System D₁C the limiting revenue rate, mean price, and throughput at each server for $\lambda < 2$ are

$$\bar{R}_{D_1C} = \frac{\lambda}{4v}, \quad \bar{P}_{D_1C} \triangleq \frac{1}{v}, \quad \rho_{D_1C} \triangleq \frac{\lambda}{2}. \quad (22)$$

The limiting revenue rate, mean price, and throughput at each server for $\lambda \geq 2$ are

$$\bar{R}_{D_1C} \triangleq \frac{(\lambda-1)}{v\lambda}, \quad \bar{P}_{D_1C} \triangleq \frac{\lambda-1}{v}, \quad \rho_{D_1C} \triangleq 1. \quad (23)$$

Proof: The result can be shown by adapting Remark 8 to the setting of *i.i.d.* exponential price. ■

Theorem 6 (R₁C). For System R₁C, the limiting revenue rate, and the mean price at each server are

$$\bar{R}_{R_1C} \triangleq \frac{1}{v} \left(\frac{\sqrt{1+\lambda}-1}{\sqrt{1+\lambda}+1} \right), \quad \bar{P}_{R_1C} \triangleq \frac{\sqrt{1+\lambda}}{v}. \quad (24)$$

The throughput at each server is

$$\rho_{R_1C} \triangleq \frac{\lambda}{\sqrt{1+\lambda}(1+\sqrt{1+\lambda})}. \quad (25)$$

Proof: Consider System R₁C with price parameter d at each server. Then, we can model the state of each server as a two state Markov chain with states $\{0,1\}$ indicating idle or busy state with thinned Poisson arrival of rate $\frac{\lambda d}{d+v}$ transitioning the server from idle to busy and unit service rate transitioning the server from busy to idle. We observe that the equilibrium probability of server being busy is $\gamma \triangleq \frac{1}{1+\frac{d+v}{\lambda d}}$, the limiting revenue rate is $\bar{\gamma} \frac{d}{(d+v)^2}$, and the acceptance probability of an incoming arrival is $\gamma \frac{d}{d+v}$. The result follows by finding the optimal price parameter that maximizes the limiting revenue rate. ■

VI. NUMERICAL RESULTS

We conduct numerical evaluations of the four approaches to illustrate their performance as the arrival rate of jobs is scaled. The metrics of interest are (i) revenue rate (ii) mean offered price and (iii) throughput of completed jobs.

Revenue Rate: We evaluate the revenue rate as a function of the job arrival rate λ at two different values of the value parameter, $v = 0.5$ and $v = 1.5$. As expected, under centralized matching and price parameter choice D₁C dominates the performance of all other systems, while simply picking one server at random in R₁C falls far short of both R₂C and the two server competition approach of R₂G. Interestingly, Fig 5 tells us that the revenue rate for the R₂C and R₂G systems appear to be nearly identical as a function of λ .




Fig. 5: Mean revenue rate as a function of λ .

We investigate this observation further using a conventional loss metric used in game theory, the *price of anarchy*. The price of anarchy measures the loss of performance due to competition when compared to collective bargaining

$$\mathcal{A} \triangleq 1 - \frac{\bar{R}_{R_2G}}{\bar{R}_{R_2C}}$$

As seen in Fig 6, there is indeed a difference in the limiting revenue rate when the arrivals are near 0, but this difference is small as λ approaches 1 and essentially disappears past $\lambda > 2$.




Fig. 6: Price of anarchy as a function of λ .

Mean Price: Fig. 7 plots the price chosen by the server for different service disciplines. Note that for $\lambda > 2$ the price for D_1C is linear while the price for R_1C is order $\sqrt{\lambda}$. This indicates that R_2C and R_2G follow a slightly higher than $\sqrt{\lambda}$ price as λ grows larger.

Normalized Throughput: The normalized throughput is shown in Fig. 8, where we see that R_2G (and R_2C) closely matches the fully centralized D_1C for arrival rate $\lambda < 1$. As the load increases, the blocking probability of randomized matching does as well, causing a reduction in throughput (which is upper-bounded by 1 in any case). Thus, for realistic traffic loads power-of-two randomized matching in the game setting performs quite effectively in all performance metrics.




Fig. 7: Tagged server price as a function of λ , (a) Tagged server price of all the systems when value parameter is set to 0.5, (b) decentralized systems for $v = 0.5$ (c) for $v = 1.5$, (d) decentralized for $v = 1.5$.




Fig. 8: Throughput as a function of λ .

VII. CONCLUSION

In this paper, we explored the dynamics of a large-scale service marketplace and proposed a distributed matching strategy that involves randomly matching jobs to two servers, each setting their own competitive prices. We demonstrated that this approach significantly mitigates the inefficiencies observed in single-server matching by ensuring lower blocking probabilities and higher system revenue. Our analysis revealed that this two-server strategy achieves revenue scaling comparable to more complex centralized and collaborative pricing models while maintaining a simpler implementation. Furthermore, we established the existence of a mean field equilibrium, providing a stable pricing environment in a competitive market setting. These findings were numerically evaluated, highlighting the practical benefits.

REFERENCES

- [1] A. A. Cournot, "Researches into the mathematical principles of the theory of wealth," *Forerunners of Realizable Values Accounting in Financial Reporting*, 1838.
- [2] I. Fisher, "Cournot and Mathematical Economics," *The Quarterly Journal of Econ.*, vol. 12, no. 2, pp. 119–138, Jan. 1898.
- [3] P. Naor, "The regulation of queue size by levying tolls," *Econometrica*, vol. 37, no. 1, pp. 15–24, Jan. 1969.

- [4] N. M. Edelson and D. K. Hilderbrand, "Congestion tolls for poisson queueing processes," *Econometrica*, vol. 43, no. 1, pp. 81–92, Jan. 1975.
- [5] D. W. Low, "Optimal dynamic pricing policies for an M/M/s queue," *Op. Res.*, vol. 22, no. 3, pp. 545–561, May-Jun 1974.
- [6] A. V. den Boer, "Dynamic pricing and learning: Historical origins, current research, and new directions," *Surveys in Op. Res. and Mgmt. Sc.*, vol. 20, no. 1, pp. 1–18, Jun. 2015.
- [7] S. Agrawal and W. Tang, "Dynamic pricing and learning with long-term reference effects," in *ACM Conf. Econom. Comput.*, vol. 25, no. 1, Dec. 2024, p. 72.
- [8] S. Agrawal, S. Yin, and A. Zeevi, "Dynamic pricing and learning under the bass model," in *ACM Conf. Econom. Comput.*, vol. 22, no. 1, Jul. 2021, pp. 2–3.
- [9] C. Borgs, J. T. Chayes, S. Doroudi, M. Harchol-Balter, and K. Xu, "The optimal admission threshold in observable queues with state dependent pricing," *Probab. in Engg. Info. Sciences*, vol. 28, no. 1, pp. 101–119, Jan. 2014.
- [10] A. K. K.S., C. K. Singh, S. T. Maguluri, and P. Parag, "Optimal pricing in multi server systems," *Perf. Eval.*, vol. 154, p. 102282, Apr. 2022.
- [11] —, "Optimal pricing in a single server system," *ACM Trans. Model. Perform. Eval. Comput. Syst.*, vol. 8, no. 4, pp. 1–32, Dec. 2023.
- [12] J. Kim and R. S. Randhawa, "The value of dynamic pricing in large queueing systems," *Op. Res.*, vol. 66, no. 2, Mar–Apr 2018.
- [13] E. A. Feinberg and F. Yang, "Optimal pricing for a $GI/M/k/N$ queue with several customer types and holding costs," vol. 82, no. 1–2, pp. 103–120, Feb. 2016.
- [14] T. X. Tran and D. Pompili, "Joint task offloading and resource allocation for multi-server mobile-edge computing networks," *IEEE Trans. Veh. Technol.*, vol. 68, no. 1, pp. 856–868, Jan. 2019.
- [15] L. Kong, J. Tan, J. Huang, G. Chen, S. Wang, X. Jin, P. Zeng, M. Khan, and S. K. Das, "Edge-computing-driven Internet of Things: A survey," vol. 55, no. 8, pp. 1–41, Aug. 2023.
- [16] J.-M. Lasry and P.-L. Lions, "Mean field games," vol. 2, pp. 229–260, Mar. 2007.
- [17] M. Huang, R. P. Malhamé, and P. E. Caines, "Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle," vol. 6, no. 3, pp. 221–252, 2006.
- [18] M. Manjrekar, V. Ramaswamy, V. R. Raja, and S. Shakkottai, "A mean field game approach to scheduling in cellular systems," *IEEE Trans. Control Netw. Syst.*, vol. 7, no. 2, pp. 568–578, Jun. 2020.
- [19] J. Li, R. Bhattacharyya, S. Paul, S. Shakkottai, and V. Subramanian, "Incentivizing sharing in realtime D2D streaming networks: A mean field game perspective," *IEEE/ACM Trans. Netw.*, vol. 25, no. 1, pp. 3–17, Feb. 2017.
- [20] R. A. Banez, H. Tembine, L. Li, C. Yang, L. Song, Z. Han, and H. V. Poor, "Mean-field-type game-based computation offloading in multi-access edge computing networks," *IEEE Trans. Wireless Commun.*, vol. 19, no. 12, pp. 8366–8381, Dec. 2020.
- [21] D. Narasimha, S. Shakkottai, and L. Ying, "A mean field game analysis of distributed MAC in ultra-dense multichannel wireless networks," *IEEE/ACM Trans. Netw.*, vol. 28, no. 5, pp. 1939–1952, Oct. 2020.
- [22] —, "Age-dependent distributed MAC for ultra-dense wireless networks," *IEEE/ACM Trans. Netw.*, vol. 31, no. 4, pp. 1674–1687, Aug. 2023.
- [23] S. Aggarwal, M. A. uz Zaman, M. Bastopcu, S. Ulukus, and T. Başar, "A mean field game model for timely computation in edge computing systems," 2024, arXiv:2404.02898 [cs.IT].
- [24] W. Mao, H. Qiu, C. Wang, H. Franke, Z. Kalbarczyk, R. K. Iyer, and T. Başar, "A mean-field game approach to cloud resource management with function approximation," in *Conf. Neural Info. Process. Sys. (NeurIPS)*, vol. 36, no. 1, 2022, pp. 36243–36258.
- [25] J. Li, B. Xia, X. Geng, H. Ming, S. Shakkottai, V. Subramanian, and L. Xie, "Mean field games in nudge systems for societal networks," *ACM Trans. Model. Perform. Eval. Comput. Syst.*, vol. 3, no. 4, pp. 1–31, Dec. 2018.
- [26] B. Xia, S. Shakkottai, and V. Subramanian, "Small-scale markets for bilateral resource trading in the sharing economy," in *IEEE Conf. Comp. Commun. (INFOCOM)*, 2018, pp. 2447–2455.
- [27] M. Mitzenmacher, "The power of two choices in randomized load balancing," *IEEE Trans. Parallel Distrib. Syst.*, vol. 12, no. 10, pp. 1094–1104, Oct. 2001.
- [28] N. D. Vvedenskaya, R. L. Dobrushin, and F. I. Karpelevich, "Queueing system with selection of the shortest of two queues: An asymptotic approach," *Prob. Info. Transmission*, vol. 32, no. 1, pp. 15–27, 1996.
- [29] L. Ying, "On the approximation error of mean-field models," in *ACM SIGMETRICS Inter. Conf. Meas. Model. Comp. Sci.*, Jun. 2016, pp. 285–297.
- [30] R. W. Wolff, "Poisson arrivals see time averages," *Op. Res.*, vol. 30, no. 2, pp. 223–231, Mar–Apr 1982.
- [31] C. Berge, *Topological Spaces: including a treatment of multi-valued functions, vector spaces and convexity*. UK: MacMillan, 1963.
- [32] L. E. J. Brouwer, "Über abbildung von mannigfaltigkeiten," *Mathematische Annalen*, vol. 71, pp. 97–115, Mar. 1911.
- [33] C. Fricker, P. Robert, and D. Tibi, "On the rates of convergence of erlang's model," *J. App. Prob.*, vol. 36, no. 4, pp. 1167–1184, 1999.