The Power of Two in Large Service-Marketplaces

Dheeraj Narasimha, Srinivas Nomula, *Student Member, IEEE*, Srinivas Shakkottai, *Senior Member, IEEE*, Parimal Parag, *Senior Member, IEEE*

***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 higher blocking, resulting in lower throughput and reduced revenue. We then examine matching jobs to two servers, which compete to provide service if unoccupied. We demonstrate the existence of 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 ensures lower blocking probabilities and higher system revenue. Our findings are validated through simulations illustrating a variety of operating scenarios.**

I. Introduction

The rapid expansion of cloud-based service ecosystems has reshaped how computational tasks are provisioned. These platforms support a wide array of workloads, such as containerized applications, artificial intelligence training and inference, and decentralized ledger operations including blockchain.

The associated service markets such as public cloud infrastructure (e.g., AWS EC2 spot/on-demand markets), GPU clusters for AI training/inference, and edge-computing providers, typically comprise a large number of servers operated by different firms competing for job requests, each catering to a continuous and diverse stream of job arrivals. The jobs correspond to workloads such as short-lived containerized tasks, ML training or inference jobs with varying latency sensitivity, or blockchain validation tasks, which differ in urgency (delay tolerance), computational intensity (CPU/GPU demand), and willingness to pay, resulting in a complex demand landscape. In such a setting, the central challenge lies in developing mechanisms that can efficiently allocate jobs to resources while maximizing revenue and minimizing service denials.

Traditional infrastructures approach this problem through centralized coordination, where a global controller assigns jobs to available servers and establishes uniform service pricing. Although such systems simplify price-setting and assignment logic, they suffer from scalability issues under high-load scenarios. Specifically, centralized schemes require real-time global knowledge of server states and continuous optimization, which become computationally infeasible as the system grows. The ongoing surge in demand for cloud services underscores the necessity for market architectures that are both adaptive and scalable.

To address these limitations, decentralized marketplaces have emerged as a compelling alternative. In these systems, jobs are routed to servers via randomized mechanisms, and each server independently determines its service price based on local observations such as current load and the statistical distribution of job valuations. This autonomy allows servers to adjust prices dynamically in response to fluctuating demand. However, decentralization introduces its own challenges, particularly in maintaining system-wide efficiency and economic viability at scale.

In this work, we explore a fundamental design question: *Is it possible to build a simple, randomized job-server matching framework in which each server autonomously sets its own price, yet the overall system still achieves high throughput and revenue?*

Main Results

We explore the design of a large-scale service marketplace by focusing on two key decisions: (i) the mechanism used to match incoming jobs to servers and (ii) the strategy for determining service prices. Under a centralized matching scheme, incoming jobs are directed to idle servers through a global gateway, assuming such servers are available. While effective in minimizing job blocking, this design necessitates a central controller that constantly monitors all server states, which introduces coordination and scaling bottlenecks. A simpler alternative is to assign jobs to servers randomly, potentially increasing the blocking probability but significantly reducing coordination complexity.

We propose a pricing model in which each server quotes a service price sampled from an exponential distribution. The rationale for choosing this distribution is twofold: it injects a small degree of randomness into the pricing to avoid pathological scenarios such as tie-breaking failures, while still concentrating values near the mean due to the distribution's light

Dheeraj Narasimha is with the Laboratoire d'Informatique de Grenoble, 38401, France. Srinivas Nomula and Parimal Parag are with the Department of Electrical Communication Engineering, Indian Institute of Science, Bangalore, Karnataka 560012, India. Srinivas Shakkottai 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.

An earlier version of this work was presented in part at IEEE INFOCOM 2025 [1].

The work of Srinivas Shakkottai is supported in part by National Science Foundation under Grant CNS 2312978. The work of Parimal Parag is supported in part by Qualcomm Inc. under Qualcomm University Relations 6G India; in part by Anusandhan National Research Foundation (ANRF) under Grant CRG/2023/008854; and in part by Centre for Networked Intelligence (a Cisco Corporate Social Responsibility (CSR) Initiative), IISc. Any opinions, findings, conclusions, and recommendations expressed in this material are those of the authors and do not necessarily reflect the views of funding agencies.

tail. Two pricing strategies are analyzed: in the first, servers cooperatively select the distribution's parameter to optimize collective revenue; in the second, each server independently optimizes its parameter in a competitive setting.

We study the following configurations across the decision rules for matching and pricing.

1. $\mathrm{D_1C}$: Jobs are deterministically matched to a vacant server by a centralized entity, and prices are collaboratively set by all servers. This idealized setup assumes complete system-level coordination and serves as a performance benchmark.
2. $\mathrm{R_2C}$: Each job is matched to two randomly chosen servers, while price parameters are still chosen cooperatively. This reduces the complexity of centralized matching but preserves coordinated pricing.
3. $\mathrm{R_2G}$: Jobs are matched to two servers at random, and each server sets prices strategically based on its beliefs about others' price distributions. This setting eliminates both centralized matching and collaborative pricing, and we model it as a mean field game (MFG), wherein servers optimize their prices given a belief about the price distribution of the rest of the population.
4. $\mathrm{R_1C}$: Each job is randomly matched to a single server, and the price parameter is selected centrally. While this approach simplifies matching, it does not incorporate pricing competition.

We perform an analytical comparison of these configurations in terms of revenue generation and system throughput. Our focus is on evaluating $\mathrm{R_2G}$, which leverages the benefits of randomized dual-server matching alongside decentralized competitive pricing. Through numerical studies, we observe that $\mathrm{R_2G}$ yields revenue and throughput close to that of $\mathrm{R_2C}$. A key technical contribution is demonstrating convergence to the mean field limit and the establishment of a consistent mean field game equilibrium.

In a mean field game formulation, a representative agent models the rest of the system as a continuum of homogeneous agents playing a fixed strategy, known as the mean field limit, and optimizes its own strategy in response. Each agent views itself as the tagged server and all others as untagged; due to symmetry, the same strategy is adopted by all servers. In our system, the prices selected by the other homogeneous servers form the mean field distribution against which the representative server has to optimize its strategy.

In practice, platforms usually estimate demand distributions and mean willingness-to-pay using posted-price acceptance rate data. For example, historical price data is directly accessible in major cloud markets (e.g., via the AWS EC2 Spot Price History API), enabling empirical calibration of demand and acceptance behavior. The pricing parameter in our scheme described next can be updated via empirical learning, without requiring exact prior knowledge of the valuation distribution. Moreover, we will show using simulations under non-exponential valuation models (Gaussian, Beta, Log-normal) that the equilibrium behavior and performance gains are robust to exact distribution specification.

Our mean field game analysis reveals that decentralized pricing in conjunction with random two-server matching mitigates the high blocking problem typically encountered with increasing arrival rates in random single server matching. This approach substantially reduces job rejection rates and enhances profitability when compared with the single-match scenario $\mathrm{R_1C}$. Despite its simplicity, it performs nearly as well as centralized and collaborative schemes.

Our methodological advance lies in proving convergence to the mean field limit in the $\mathrm{R_2G}$ regime. Unlike standard proofs that assume identical policies across agents, we explicitly model a "tagged server" as an agent who strategically responds to the price distribution of others. This framework allows us to explicitly define the game between the tagged agent and the rest of the ensemble. We note that while our methods also apply to games such as $\mathrm{D_1G}$, which consists of matching jobs to a vacant server that implicitly participates in a price competition, the performance would naturally be upper bounded by $\mathrm{D_1C}$. Our emphasis in this work is on the benefits of multiple randomized matches, and so we do not pursue details of such other scenarios here.

Finally, we corroborate our theoretical insights with numerical experiments across a range of settings. These simulations demonstrate that modest increases in the number of server choices can lead to significant gains in system efficiency and profitability. Through numerical studies, we also observe that the insights obtained from the analytical study for exponential valuation and price distributions continue to hold true for a number of other price and value distributions as well. Our results indicate a promising direction for designing scalable, decentralized service markets with minimal coordination overhead.

Related Works

While dynamic pricing has been studied since the early work of Cournot [2], further elaborated in [3], the use of pricing as a control mechanism can be traced back to [4]–[6]. For a more comprehensive overview of dynamic pricing, we refer the reader to [7]–[9], among others.

More recently, there has been a growing focus on dynamic pricing in the context of cloud computing [10]–[13], as well as in multi-class, multi-server models such as [14]. These models typically adopt the viewpoint of a central manager operating with a limited pool of servers.

In contrast, our work departs from this centralized perspective and instead views the problem as a game played by independent agents; see, for example, [15], [16] and the references therein. This decentralized viewpoint becomes analytically challenging, as the revenue rate for each agent depends on the pricing strategies of all others. To address this, we adopt a mean field game formulation. In this framework, the system is viewed from the perspective of a representative agent interacting with a continuum of others.

The mean field game approach originated independently in the works of [17] and [18]. This perspective has since been applied to various networked systems [19]–[25] and networked markets [26], [27]. The former focus on competition in resource-constrained environments, while the latter align more closely with our interest in market design. However, these prior works do not explicitly address service markets with server-level competition. For instance, [27] studies one-to-one matching of consumers to producers in a prosumer market, and thus does not encounter occupancy-related blocking or leverage the power of two choices, as our model does.

As with these related works, our objective is to harness the analytical tractability of the mean field approach while enabling the numerical computation of fixed points corresponding to control policies in complex systems. This allows us to obtain strong convergence results that are of independent interest, while directly addressing the challenges of server matching and pricing in compute service markets.

We consider a setting with homogeneous servers and unit service capacity in order to understand the interaction between randomized matching and decentralized pricing. Prior works such as [28]–[30] study more general heterogeneous loss systems and packing constraints, with a focus on allocation efficiency and stability. In contrast, our work focuses on strategic pricing and equilibrium behavior in decentralized markets, which introduces additional analytical challenges. Even in this simplified setting, establishing mean field convergence and equilibrium requires careful analysis. The present model can thus be viewed as a tractable baseline that captures the key effects of competition and limited choice.

Notation: We denote the set of first N integers by $[N] \triangleq \{1,\dots,N\}$, the set of all positive integers by $\mathbb{N}$, the set of all real numbers by $\mathbb{R}$, the set of nonnegative real numbers by $\mathbb{R}_+$, the set of all functions from a domain A to a co-domain B by B^A, and $\bar{z} \triangleq 1-z$ for all $z \in [0,1]$.

II. System Model

We consider a system with $N+1$ servers labeled $\mathcal{N} \triangleq \{0,\dots,N\}$. Tasks arrive in the system according to a Poisson process with homogeneous rate $(N+1)\lambda$. Let A_k denote the arrival time of the kth task, and $T_k \triangleq A_k - A_{k-1}$ its inter-arrival time, for all $k \in \mathbb{N}$. Thus, the task arrival time sequence is represented by $A \in \mathbb{R}_+^{\mathbb{N}}$, and the inter-arrival time sequence by $T \in \mathbb{R}_+^{\mathbb{N}}$, where T is an *i.i.d.* sequence of exponential random variables with rate $(N+1)\lambda$.

We assume that the service time for any task at any of the $N+1$ servers is an *i.i.d.* exponentially distributed random variable with rate 1. Let $S \in \mathbb{R}_+^{\mathbb{N}}$ denote the *i.i.d.* sequence of service times, where S_k is the service requirement of task k arriving at time A_k.

Each server can be in one of two states, $\{0,1\}$, representing *idle* and *busy*, respectively. A server transitions from idle to busy upon accepting a task, and from busy to idle upon completing service. The state space of all servers is denoted by $\mathcal{X} \triangleq \{0,1\}^{\mathcal{N}}$, where the state vector $x \in \mathcal{X}$ specifies the occupancy of all servers, and $x_n \in \{0,1\}$ indicates the state of server $n \in \mathcal{N}$. The overall system state is described by the process $X \in \mathcal{X}^{\mathbb{R}_+}$, where $X_t \in \mathcal{X}$ denotes the system state at time $t \in \mathbb{R}_+$. The random variable $X_{t,n} \in \{0,1\}$ indicates the occupancy of server $n \in \mathcal{N}$ at time t.

Server 0 is referred to as the tagged server, while the remaining N servers in $[N] = \mathcal{N} \setminus \{0\}$ are considered untagged. We define the fraction of occupied untagged servers $\{1,\dots,N\}$ at time t by

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

As discussed earlier, in a mean field game, a representative agent views the rest of the system as a large population of identical agents following a fixed strategy, referred to as the mean field limit, and then chooses its own strategy in response. Each agent treats itself as the tagged server and all other servers as untagged. Because of symmetry, the same strategy is used by all servers.

A. Task Valuation and Server Pricing

Let $V \in \mathbb{R}_+^{\mathbb{N}}$ denote the sequence of random valuations for incoming tasks, where task k has a positive valuation $V_k \in \mathbb{R}_+$. We assume that V is an *i.i.d.* sequence with common distribution function $G : \mathbb{R}_+ \to [0,1]$ and complementary distribution $\bar{G} \triangleq 1 - G$. We assume the valuation distribution is exponential with unit rate.[1]

Each server n independently sets a random price $P_{k,n} : \Omega \to \mathbb{R}_+$ for each incoming task k. The tagged server 0 is modeled as a rational agent, while all other servers follow a common random pricing strategy. Let $F_0 : \mathbb{R}_+ \to [0,1]$ denote the price distribution at the tagged server 0, and $F_1 : \mathbb{R}_+ \to [0,1]$ the price distribution at each untagged server $n \in [N]$.

We consider the case where both F_0 and F_1 are exponential distributions with rates d_0 and d_1, respectively:

$$F_i(x) \triangleq 1 - e^{-d_i x}, \quad \text{for all } x \in \mathbb{R}_+ \text{ and } i \in \{0,1\}.$$

B. Server Selection and Joining

Our approach to selecting servers at random is inspired by a widely-used load-balancing technique known as the power-of-two choices [31], [32]. We adapt this method to suit our setting as follows.

Each task k arrives at a dispatcher that selects a pair of distinct servers $I_k \subseteq \mathcal{N}$ uniformly at random without replacement. If both servers in I_k are currently busy, task k departs the system without being served. If one of the servers $n \in I_k$ is idle, the task queries the server for its asking price. If the task's valuation V_k exceeds the server's price $P_{k,n}$, it joins the idle server; otherwise, it exits the system.

If both servers in I_k are idle, they compete for the task. The task joins the server offering the lower price, provided its valuation exceeds this price, i.e., $V_k > \min\{P_{k,n} : n \in I_k\}$. If both servers quote the same price, the task chooses between them uniformly at random, joining either with probability $\frac{1}{2}$.

For a system with $(N+1)$ servers, we define ξ_k^N as the indicator that task k joins the tagged server 0. Using this indicator, we express the average revenue rate earned by the tagged server 0 over the interval $[0,t]$ as

$$R_t^N \triangleq \frac{1}{t}\sum_{k=1}^{K_t} P_{k,0}\xi_k^N, \tag{1}$$

[1] All results for a general valuation rate can be recovered via scaling with the valuation rate.

where $K_t \triangleq \sum_{n\in\mathbb{N}} \mathbb{1}_{\{A_n \leqslant t\}}$ is the total number of arrivals up to time t.

The key performance metric is the limiting revenue rate as the system size becomes large:

$$R \triangleq \lim_{N\to\infty} \lim_{t\to\infty} R_t^N. \tag{2}$$

Our motivation for choosing exponential pricing arises from the analytical complexity introduced by more general pricing distributions. Even relatively simple parameterized distributions, such as the Gaussian, exhibit complicated behaviors when considering the minimum over multiple random draws, especially when analyzing best-response strategies. Nevertheless, we provide simulation results that demonstrate the effectiveness of our approach under broader pricing and valuation distributions.

III. Mean field analysis

We analyze the joint evolution of the process $(X_{t,0}, Z_t^N, t \geqslant 0)$, where $X_{t,0}$ denotes the occupancy status of the tagged server 0 at time $t \in \mathbb{R}_+$, and Z_t^N is the empirical fraction of occupied untagged servers in $[N]$. We denote the history of the process up to time t as $\mathcal{F}_t^N \triangleq \sigma(X_{s,n}, n \in \mathcal{N}, s \leqslant t)$ for all $N+1$ servers, and the history up to the kth arrival time A_k as $\mathcal{F}_{A_k}^N$.

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

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

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

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

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

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

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

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

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

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

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 unit rate, we have*

$$q_1 = \frac{d_0}{d_0+1}, \quad q_{20} = \frac{d_0}{d_0+d_1+1}, \quad q_{21} = \frac{d_1}{d_0+d_1+1},$$
$$p_1 = \frac{d_1}{d_1+1}, \quad p_2 = \frac{2d_1}{2d_1+1}.$$

Proof: The result follows from the distribution of minimum of independent exponentially distributed random variables. For details, please see Appendix A-A. ■

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

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

Proof: Please see Appendix A-B. ■

Recall that a randomly selected subset $I_k \in \binom{\mathcal{N}}{2}$, 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 $j \in \{0,1\}, n \in [N]$, as

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

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

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

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

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*

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

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: The result follows from uniform sampling of two server subsets of $(N+1)$ server set $\mathcal{N}$. For details, please see Appendix A-C. ■

A. System evolution

We will show that the joint evolution of busyness of the tagged server and the empirical fraction of busy untagged servers is Markov, and characterize the associated generator matrix.

Proposition 1. *The process $(X_{t,0}, Z_t^N, t \geqslant 0)$ is a continuous time Markov chain with the following properties.*

(a) ***Generator matrix*** *The associated generator matrix Q^N for each $(x,z) \neq (y,w) \in \{0,1\} \in \mathcal{Z}_N$ is*

$$Q^N_{(x,z),(x,z-\frac{1}{N})} = Nz,$$
$$Q^N_{(x,z),(x,z+\frac{1}{N})} = \lambda\bar{z}[2p_1(x+Nz) + 2\bar{x}q_{21} + p_2(N\bar{z}-1)],$$
$$Q^N_{(1,z),(0,z)} = 1,$$
$$Q^N_{(0,z),(1,z)} = 2\lambda(zq_1 + \bar{z}q_{20}).$$

(b) ***Positive recurrence.*** *If $d_0, d_1 \in (0,\infty)$, then the process is positive recurrent with stationary distribution π^N.*

Proof: Transition rate results follows from the exponential holding rates for each state of the joint process, and positive recurrence follows from the irreducibility of finite state embedded discrete time Markov chain and finite transition rates. For details, please see Appendix A-D. ∎

Definition 3. For CTMC $(X_{t,0}, Z^N_t, t \geqslant 0)$, we define the mean rate of change of fraction of busy untagged servers Z^N_t as a map $h_N : \{0,1\}\times\mathcal{Z}_N \to \mathbb{R}$ for each $(x,z) \in \{0,1\}\times\mathcal{Z}_N$ as

$$h_N(x,z) \triangleq \sum_{(y,w)\in\{0,1\}\times\mathcal{Z}_N} Q^N_{(x,z),(y,w)}(w-z). \tag{3}$$

Lemma 4. *The rate of change of fraction of busy untagged servers for each $(x,z) \in \{0,1\}\times\mathbb{R}$ is*

$$h_N(x,z) = \lambda\bar{z}[2p_1(\frac{x}{N}+z) + 2\frac{\bar{x}}{N}q_{21} + p_2(\bar{z}-\frac{1}{N})] - z.$$

Proof: The result follows from the definition of h_N in Definition 3 and the form of generator matrix Q^N in Proposition 1(a). ∎

B. Mckean-Vlasov equation

We observed that the empirical fraction of busy untagged servers is Markov. In this section, we will show that under certain conditions this *stochastic* evolution converges to a *deterministic* evolution for large N. This deterministic evolution is characterized by Mckean-Vlasov equation. For the proposed system, this equation is defined in Definition 5.

Definition 4. Consider a family of CTMCs with generator matrix Q^N for each $N \in \mathbb{N}$ and the associated map $h_N : \{0,1\}\times[0,1] \to \mathbb{R}$ in Definition 3. For this family of CTMCs we define the point-wise limit $h \triangleq \lim_{N\to\infty} h_N$ and call it *mean field model*.

Lemma 5. *For the family of CTMCs with generator matrix Q^N in Proposition 1, the associated limiting map $h : [0,1] \to \mathbb{R}$ doesn't depend on $x \in \{0,1\}$ and is defined as*

$$h(z) = \lambda\bar{z}(2zp_1 + \bar{z}p_2) - z, \quad z \in [0,1]. \tag{4}$$

Proof: From the form of h_N and generator matrix Q^N in Lemma 4 and Proposition 1 respectively, we obtain

$$h_N(x,z) = \frac{\lambda\bar{z}}{N}\Big[(2p_1-p_2)x + (2q_{21}-p_2)\bar{x}\Big] + h(z). \tag{5}$$

The result follows since $\lim_{N\to\infty}|h_N(x,z)-h(z)| = 0$ and hence the point-wise limit $\lim_{N\to\infty} h_N$ doesn't depend on x. ∎

Definition 5. [Mckean-Vlasov equation] In terms of the map $h : [0,1] \to \mathbb{R}$ defined in (4), we can define an autonomous nonlinear system $\Phi : \mathbb{R}_+\times[0,1] \to [0,1]$ for any time $t \in \mathbb{R}_+$ and initial state $z \in [0,1]$ as

$$\Phi_t(z) \triangleq z + \int_0^t h(\Phi_s(z))ds, \qquad \Phi_0(z) = z. \tag{6}$$

This equation is referred to as *Mckean-Vlasov equation* and determines the evolution of the deterministic mean field model. For this autonomous nonlinear system,we define the set of rest points as

$$\mathcal{S} \triangleq \{z \in [0,1] : h(z) = 0\}.$$

Definition 6. We define the the evolution of the distance of Mckean-Vlasov equation from a rest point $z^* \in \mathcal{S}$ as $\varepsilon : \mathbb{R}_+\times[0,1] \to \mathbb{R}$ for any initial point $z \in [0,1]$ at time $t \in \mathbb{R}_+$, as

$$\varepsilon(t,z) \triangleq \Phi_t(z) - z^*. \tag{7}$$

Definition 7. For convenience of notation, we define the following two positive constants

$$L_\lambda \triangleq 2\lambda(2p_1-p_2) = \frac{4\lambda d_1^2}{(1+d_1)(1+2d_1)}, \tag{8}$$
$$K_\lambda \triangleq 1 + 2\lambda(p_2-p_1) = 1 + \frac{2\lambda d_1}{(1+d_1)(1+2d_1)}. \tag{9}$$

Remark 2. From the definition of constants K_λ, L_λ in Definition 7, we observe that

$$\frac{K_\lambda-1}{2\lambda} = \Big(\frac{1}{1+d_1} - \frac{1}{1+2d_1}\Big),$$
$$\frac{L_\lambda}{2\lambda} = \Big(1 - \frac{2}{1+d_1} + \frac{1}{1+2d_1}\Big) = \frac{d_1}{1+d_1} - \frac{(K_\lambda-1)}{2\lambda}.$$

This relationship implies that $K_\lambda \geqslant 1$ and the ratios $\frac{L_\lambda}{2\lambda}$ and $\frac{K_\lambda-1}{2\lambda}$ depend only on the revenue rate d_1 and do not depend on normalized arrival rate λ. Further, $\frac{L_\lambda+K_\lambda-1}{2\lambda} \in [0,1]$ is increasing in revenue rate d_1.

Theorem 1. *Consider the family of CTMCs $((X_{t,0}, Z^N_t, t \geqslant 0) : N \in \mathbb{N})$ with the associated generator matrices as defined in Proposition 1, and the associated autonomous nonlinear map Φ evolving under map $h : [0,1] \to \mathbb{R}$ defined in* (4)*. If $d_0, d_1 \in (0,\infty)$ and normalized arrival rate λ is finite, then the following statements are true.*

(a) ***Asymptotically accurate mean field model.*** *The map $h(z)$ is an asymptotically accurate mean field model. That is,*

$$\sup_{(x,z)\in\{0,1\}\times\mathcal{Z}_N} |h(z) - h_N(x,z)| \leqslant \frac{\lambda}{N}. \tag{10}$$

(b) ***Lipschitz partial derivatives.*** *The derivative $h'(z)$ exists and is affine. That is, $h'(z) = -K_\lambda - L_\lambda z$ for each $z \in [0,1]$. Hence h' is L_λ Lipschitz and uniformly bounded as*

$$K_\lambda \leqslant |h'(z)| \leqslant K_\lambda + L_\lambda, \quad z \in [0,1]. \tag{11}$$

(c) Autonomous nonlinear system Φ has a unique rest point

$$z^* \triangleq \frac{1}{L_\lambda}\Big(-K_\lambda + \sqrt{K_\lambda^2 + L_\lambda^2 + 2L_\lambda(K_\lambda-1)}\Big). \tag{12}$$

(d) ***Global exponential stability.*** *Unique rest point z^* is globally exponentially stable. In particular, for any initial condition $z \in [0,1]$ and time $t \in \mathbb{R}_+$,*

$$|\varepsilon(t,z)| \leqslant |\varepsilon(0,z)|\, e^{-\lambda p_2 t}. \tag{13}$$

(e) ***Bounded mean transition-rate condition.*** *Mean rate of change in empirical faction of busy untagged servers has a bounded second moment, i.e.*

$$\sum_{(y,w)} Q^N_{(x,z),(y,w)} |w-z|^2 \leqslant \frac{C_\lambda}{N} + \frac{\lambda}{N^2}, \tag{14}$$

where finite positive constant $C_\lambda \triangleq \sup_{z\in[0,1]}(h(z)+2z)$.

Proof: The results follow from the specific form of mean-field model h given in (4). For details, please see Appendix A-E. ■

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 (12) of the deterministic evolution governed by Mckean-Vlasov equations, together with the associated convergence rate. A visual representation of this convergence for a specific numerical example is plotted in Fig. 1.

Theorem 2. *If the conditions of Theorem 1 hold and N is large enough such that $NK_\lambda > L_\lambda$, then the stationary fraction of occupied tagged servers Z^N_∞ converges in the mean-square sense with rate $1/N$, i.e.*

$$\mathbb{E}\left|Z^N_\infty - z^*\right|^2 \leqslant \frac{2\lambda}{NK_\lambda} + \Big(\frac{1}{2K_\lambda} + \frac{3L_\lambda}{K^2_\lambda - \frac{1}{N^2}L^2_\lambda}\Big)\Big(\frac{\lambda}{N^2} + \frac{C_\lambda}{N}\Big).$$

Proof: From Theorem 1, it follows that the conditions (a),(b),(d), and (e) hold. It can then be shown that the convergence result holds under these conditions. For details, please see Appendix A-G. ■

Remark 3. Typically to prove theorems similar to Theorem 2, we have condition (e) 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. Please see [33] for a full elaboration on these conditions.

Before we proceed further to pricing, we illustrate the time evolution of the squared difference $(Z^N_t - z^*)^2$ for $(N+1)$ server systems in Fig. 1 where $N \in \{10, 100, 1000\}$, normalized task arrival rate $\lambda = 5$ each having *i.i.d.* exponentially distributed valuation with unit rate, and the price rate for the tagged server is $d_0 = 5$ and for any untagged server is $d_1 = 9$. We note that as N increases, the empirical mean occupancy Z^N_t moves towards the rest point $z^* = 0.895$ computed numerically from (12).

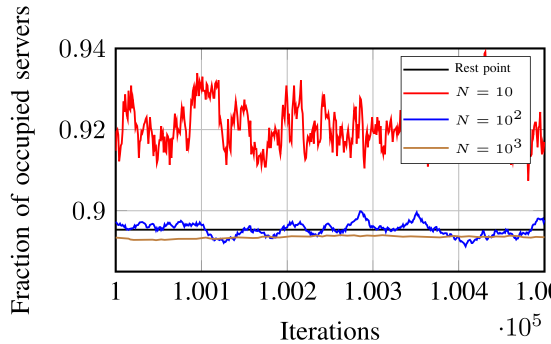

Fig. 1: Convergence of Z^N_t to the fixed point z^*.

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^N_t of occupied untagged servers $[N]$. We have shown that the stationary fraction of occupied untagged servers converges in mean square sense to z^* for large N. In the following theorem, we show that the joint distribution of $(X_{\infty,0}, Z^N_\infty)$ 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 being dependent on the deterministic limit z^*.

Theorem 3. *Consider the CTMC $(X_{t,0}, Z^N_t : t \geqslant 0)$ defined in Proposition 1 with associated stationary distribution π^N. Let $z^* \in [0,1]$ be the unique rest point of the associated mean field model (6). Then stationary distribution π^N converges pointwise to $\pi \in \mathcal{M}(\{0,1\} \times [0,1])$ in N, where for each $(x,z) \in \{0,1\} \times [0,1]$ and $\bar{z} \triangleq 1 - z$,*

$$\pi_{x,z} \triangleq \lim_{N\to\infty} \pi^N_{x,z} = \mathbb{1}_{\{z=z^*\}} \frac{\bar{x} + 2x\lambda(zq_1 + \bar{z}q_{20})}{2\lambda(zq_1 + \bar{z}q_{20}) + 1}. \tag{15}$$

Proof: The proof uses the fast simulation results from [34], [35] to show that $\pi^N_{x|z} \to \pi_{x|z^*}$ for discrete time Markov processes. Noting that $\pi^N_z \to \pi_{z^*}$ from Theorem 2 completes the proof. Please see Appendix A-H for details. ■

Remark 4. Note, this result indicates that one can treat the tagged server as largely independent of any individual server and instead approximate all interactions with the untagged server through a deterministic value, namely z^*. This is a common, yet important feature of mean field models that allows us to greatly simplify our analysis and motivates this choice of modeling our system.

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 unit rate. 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 the 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 8. For a fixed unit value rate, the mean field game equilibrium is defined by the pairwise fixed point in empirical measure and price parameter (z_M, d_M), given by

$$z_M = z^*(d_M), \qquad d_M = d_0(z_M, d_M). \tag{16}$$

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

$$R = \lim_{N\to\infty}\lim_{t\to\infty} R_t^N = \frac{z^* q_1^2 + \bar{z}^* q_{20}^2}{d_0(\frac{1}{2\lambda} + z^* q_1 + \bar{z}^* q_{20})}. \tag{17}$$

Proof: Consider the positive recurrent CTMC $(X_{t,0}, Z_t^N, t \geqslant 0)$ with generator matrix Q^N defined in Proposition 1(a), and associated invariant distribution π^N. We can write the time average of the revenue rate in a finite time duration $[0,t]$ of K_t arrivals as

$$R_t^N = \frac{K_t}{t}\Big(\frac{1}{K_t}\sum_{k=1}^{K_t} P_{k,0}\xi_k^N\Big).$$

From the strong law of large numbers, it follows that $\lim_{t\to\infty}\frac{K_t}{t} = (N+1)\lambda$ 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 kth incoming task, we can write its conditional mean $\mathbb{E}[\xi_k^N \mid X_{A_k,0}, Z_{A_k}^N, P_{k,0}]$ given the states $X_{A_k,0}, Z_{A_k}^N$ and price $P_{k,0}$ at the tagged server, as

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

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

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

First ,we recall that $K_t \to \infty$ as $t \to \infty$ for $\lambda > 0$. Second, we know that Poisson arrival sees time averages, from the PASTA [36] property. Therefore,

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

The result follows from taking limit $N \to \infty$ and using the expression for invariant distribution π from Theorem 3. ■

Remark 5. 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 for the tagged server 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. 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 $(N+1)$ server systems in Fig. 2, where $N \in \{10, 10^2, 10^2\}$, normalized task arrival rate $\lambda = 5$, where each task has an *i.i.d.* valuation distributed exponentially with unit rate, and the price rate for the tagged server is $d_0 = 5$ and for any untagged server $d_1 = 9$. We note that as N increases, the empirical revenue rate $R_{N,t}$ moves towards the limiting mean $R = 0.14$ computed numerically from (17) where $z^* = 0.895$ computed from (12).

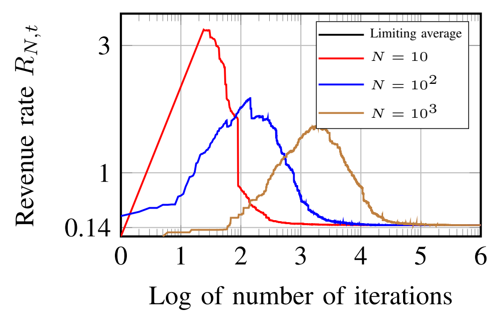

Fig. 2: Convergence of revenue rate $R_{N,t}$ to its limit R.

B. Unimodality

For fixed normalized arrival rate λ, unit rate value, and price parameter d_1 for untagged servers, let $R(d_0)$ 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 (17) to characterize the shape of $R(d_0)$.

Theorem 5. *For fixed d_1, λ, and under the large server limit, the limiting revenue rate for the tagged server is unimodal for its price parameter $d_0 \in (0, \frac{1}{\sqrt{3}-1}]$, and hence there is a unique best response*

$$d_0^* \triangleq \arg\max\Big\{R(d_0) : d_0 \in \Big(0, \frac{1}{\sqrt{3}-1}\Big]\Big\}. \tag{18}$$

Further, $R(d_0)$ is monotone decreasing for $d_0 \geqslant 1 + d_1$.

Proof: We note that a unimodal function over an interval has a unique maxima in this interval. Therefore, it suffices shows that the mean revenue rate is increasing and then decreasing in this interval. Please see Appendix B-B for details. ■

Remark 6. We note that restricting optimal price parameter $d_0^* \leqslant \frac{1}{\sqrt{3}-1}$ might seem unrealistic. However, the second part of our theorem shows that a maximizer of $R(d_0)$ can only lie in $(0,(d_1+1)]$ as seen in Lemma 9 which partially justifies this restriction. Further, we conducted numerical experiments for normalized arrival rate $\lambda = 0.5$, untagged servers' price rate $d_1 \in \{.01, .1, 1, 10\}$, 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 $d_0 \in (\frac{1}{\sqrt{3}-1}, (d_1+1)]$. One of the consequences of Theorem 5 is that there is a unique best response price parameter d_0^* for the tagged server.

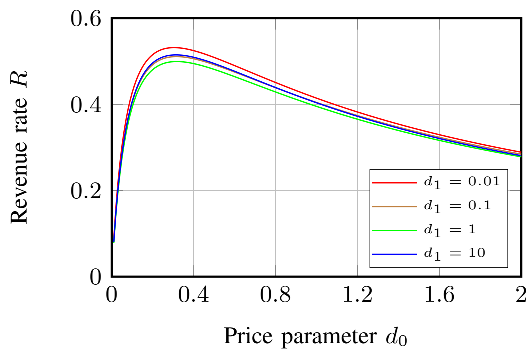

Fig. 3: Unimodality of revenue rate R as a function of tagged server price rate d_0.

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 6. *For the $(N+1)$ server system under consideration, we let the tagged server 0 select its best response price parameter d_0^* defined in (18) 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 8.*

Proof: Consider the unit value rate, fixed normalized arrival rate λ, and fixed price parameter d_1 chosen by all N untagged servers. In this setting, we know from Theorem 2 that there exists a unique fraction of busy servers $z^*(d_1)$ under large N limit at stationarity. This map $d_1 \mapsto z^*$ is continuous from Lemma 5. Next, it follows from Theorem 5 that there is a unique best response price parameter $d_0^* \in [0, \frac{1}{\sqrt{3}-1}]$ for the tagged server. Further, it can be verified that $R(d_0)$ is continuous in d_0 and differentiable for $d_0 > 0$. Hence, given z^* and d_1, from Berge's Maximum Theorem [37] it follows that d_0^* is a continuous function of (z^*, d_1). Therefore, the map $z^* \mapsto d_0^*(z^*, d_1)$ is continuous. Finally, since all agents will imitate the tagged agent's strategy, the new price rate is given by d_0^*. We summarize this relationship by

$$d_1 \to z^*(d_1) \to d_0^*(z^*(d_1), d_1) \to d_1^*.$$

Since, d_0^* always lies in the interval $[0, \frac{1}{\sqrt{3}-1}]$ which is convex and compact, we may now apply Brower's fixed point theorem [38] to the relationship $d_1 \to z^*(d_1) \to d_0^*(z^*(d_1), d_1)$ to ensure the existence of a fixed point. Hence, we obtain the desired result. ∎

Remark 7. The mean field limit z^* is computable in closed-form and given in Lemma 5. 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 unit value rate and normalized arrival rate λ.

Remark 8. In general, it is not always clear that the limit of Nash equilibrium for a game of finite N agents, converges to the mean field game equilibrium as N tends to infinity. In fact, sometimes Nash equilibria that exist in finite models disappear in the infinite model as shown in [39]. However, we note that in our homogeneous system with finite state space, it can be shown that the mean field equilibrium, (z_M, d_M) is an ϵ-Nash equilibrium with ϵ of order $O(N^{-\frac{1}{3}})$ as shown in [22, Theorem 4].

V. Comparison with other pricing policies

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

Proposition 2. *For a fixed normalized arrival rate λ and common price parameter $d_0 = d_1 = d$, we can write the limiting occupancy of $(N+1)$ server system as $z(d)$, probabilities $q_1(d) = \frac{d}{d+1}, q_{20}(d) = \frac{d}{2d+1}$. For this system, the limiting revenue rate is*

$$R(d) \triangleq \frac{z(d)q_1^2(d) + (1-z(d))q_{20}^2(d)}{d(\frac{1}{2\lambda} + z(d)q_1(d) + (1-z(d))q_{20}(d))}, \tag{19}$$

and the blocking probability of an incoming arrival is

$$p_b(d) \triangleq \Big(z^2(d) + \frac{2z(d)\bar{z}(d)}{1+d} + \frac{\bar{z}^2(d)}{1+2d}\Big)^2. \tag{20}$$

Proof: Let the price rates be $d_0 = d_1 = d$, then the selection probabilities follow. We can write the limiting fraction of occupied servers as $z(d) = z^*$ defined in (12), substituting

$$\frac{K_\lambda - 1}{2\lambda} = \frac{d}{(d+1)(2d+1)}, \quad \frac{L_\lambda}{2\lambda} = \frac{d}{d+1} - \frac{K_\lambda - 1}{2\lambda}.$$

Substituting these in the limiting revenue rate (17) at the tagged server when the price rates are equal, we get the result in (19). To obtain the blocking probability in (20), we make the following observations. First, we observe that under the mean field limit, any two sampled servers have independent occupancies and their limiting occupancy converges to $z(d)$.

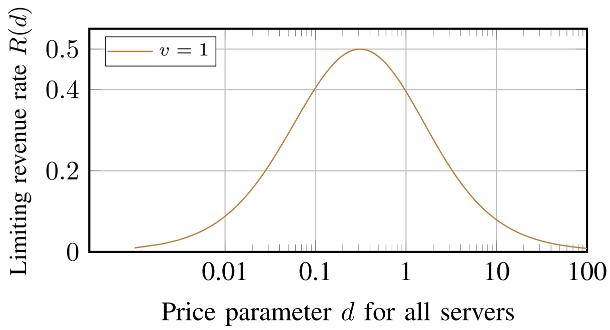

Fig. 4: Unimodality of limiting revenue rate as a function of common price parameter d for normalized arrival rate $\lambda = 5$.

Next, we observe that this system is blocked when (a) either both sampled servers are occupied, (b) or exactly one of the two sampled servers is idle and has a price rate higher than the valuation, (c) or both sampled servers are idle and the minimum price of the two idle servers is higher than the valuation. ■

Remark 9. We do not know whether the map $d \mapsto R$ 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. Part of the difficulty of the $\mathrm{R_2C}$ system comes from treating the mean field as a function of d while maximizing the revenue rate, whereas in the $\mathrm{R_2G}$ setting, the tagged server treats the limiting occupancy z^* as a constant value.

At mean field game equilibrium (z_M, d_M) of Definition 8, 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 ($\mathrm{R_2G}$)**.** *Consider System $\mathrm{R_2G}$ at the mean field game equilibrium (z_M, d_M) in Definition 8. In terms of this pair and the limiting revenue rate $R(d)$ from Proposition 2, the limiting revenue rate, the mean price, and throughput at each server are*

$$R_{\mathrm{R_2G}} \triangleq R(d_M), \quad \bar{P}_{\mathrm{R_2G}} \triangleq \frac{1}{d_M}, \quad \rho_{\mathrm{R_2G}} \triangleq \lambda \bar{p}_b(d_M). \tag{21}$$

System $\mathrm{R_2C}$ differs from System $\mathrm{R_2G}$ in that the identical price parameter at all servers is chosen to maximize the revenue rate $R_{\mathrm{R_2C}}$.

Corollary 2 ($\mathrm{R_2C}$)**.** *Consider System $\mathrm{R_2C}$ with the following unique maximizer*

$$d^* \triangleq \arg\max \{R(d) : d \in \mathbb{R}_+\}. \tag{22}$$

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

$$R_{\mathrm{R_2C}} \triangleq R(d^*), \quad \bar{P}_{\mathrm{R_2C}} \triangleq \frac{1}{d^*}, \quad \rho_{\mathrm{R_2C}} \triangleq \lambda \bar{p}_b(d^*). \tag{23}$$

Remark 10. One can draw parallels between our own exponentially priced system $\mathrm{D_1C}$ and the constant uniform price system from [11]. Concretely, consider a variant of $\mathrm{D_1C}$ system with constant uniform price p for all incoming arrivals. It follows that the thinned Poisson arrival rate to the system is $(N+1)\lambda\bar{G}(p)$. For normalized arrival rate λ and price p such that $\lambda\bar{G}(p) < 1$, it follows from [11], [40], [41] that under a large server limit with centralized routing to idle servers, the following uniform price maximizes the limiting revenue rate

$$p_U \triangleq \arg\max_p p\bar{G}(p). \tag{24}$$

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

$$p_U \triangleq \arg\max_p \left\{pe^{-p} : e \leqslant \lambda < e^p\right\} = \ln \lambda.$$

We note that the only difference between the limiting revenue rate maximizing system and $\mathrm{D_1C}$ system are the constant price in the first system and *i.i.d.* exponentially random pricing in the second system. The qualitative behavior of $\mathrm{D_1C}$ system 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 $\mathrm{D_1C}$ system, to be consistent with random exponential pricing considered in all comparison systems.

Remark 11. In the setting where servers offer constant prices and customer valuations are drawn randomly according to a distribution $G(\cdot)$, let d_0 and d_1 be two offered prices and λ denote the normalized customer arrival rate. Following a similar line of analysis, it is straightforward to observe that the long-run average revenue rate of the tagged server is

$$R(d_0) = \frac{2\lambda\bar{G}(d_0)(z + \bar{z}\mathbb{1}_{\{d_0<d_1\}})}{1 + 2\lambda\bar{G}(d_0)(z + \bar{z}\mathbb{1}_{\{d_0<d_1\}})}.$$

Here z is the mean field limit obtained when all the untagged servers choose the price d_1. This revenue function $\bar{R}(d_0)$ is generally discontinuous in d_0. Hence, it is not straight-forward to establish closed form expressions for the tagged server pricing. Therefore, we avoid explicitly analyzing this system. One can however use our methodology from the following sections to obtain approximate equilibrium results.

Proposition 3 ($\mathrm{D_1C}$)**.** *For $\mathrm{D_1C}$ system, the limiting revenue rate, mean price, and throughput at each server for $\lambda < 2$ are*

$$R_{\mathrm{D_1C}} = \frac{\lambda}{4}, \qquad \bar{P}_{\mathrm{D_1C}} \triangleq 1, \qquad \rho_{\mathrm{D_1C}} \triangleq \frac{\lambda}{2}. \tag{25}$$

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

$$R_{\mathrm{D_1C}} \triangleq \frac{(\lambda-1)}{\lambda}, \quad \bar{P}_{\mathrm{D_1C}} \triangleq \lambda - 1, \quad \rho_{\mathrm{D_1C}} \triangleq 1. \tag{26}$$

Proof: The result can be shown by adapting Remark 10 to the setting of *i.i.d.* exponential price. Please see Appendix C-A. ■

Theorem 7 ($\mathrm{R_1C}$)**.** *For $\mathrm{R_1C}$ system, the limiting revenue rate, and the mean price at each server are*

$$R_{\mathrm{R_1C}} \triangleq \left(\frac{\sqrt{1+\lambda}-1}{\sqrt{1+\lambda}+1}\right), \qquad \bar{P}_{\mathrm{R_1C}} \triangleq \sqrt{1+\lambda}. \tag{27}$$

The throughput at each server is

$$\rho_{R_1C} \triangleq \frac{\lambda}{\sqrt{1+\lambda}(1+\sqrt{1+\lambda})}. \tag{28}$$

Proof: The computations are similar to the ones performed in Theorem 4. Please see Appendix C-B for details. ∎

VI. Numerical and Simulation Results

In this section, we evaluate the performance of our system and present key insights. Our numerical results demonstrate that protocols derived from mean field analysis perform well, even in low-information settings. We then present numerical methods to compute the mean-field game equilibrium. Finally, we provide simulation-based studies that extend our results to more general price and valuation models. Examples include Log-normal, Gaussian, and Beta distributions. Our simulations suggest that insights derived from the analytically tractable exponential case are useful for computing fixed points and evaluating performance in these more general settings.

A. Numerical Results

We conduct numerical evaluations for the four systems of interest to illustrate their performance as the job arrival rate is scaled with the system size. 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 normalized job arrival rate λ for the value parameter $v = 1$. As expected, under centralized matching with revenue maximizing price parameter, D_1C has the highest revenue rate when compared to the three other systems, uniformly for all normalized arrival rates λ. In contrast, selecting a single server at random as in R_1C, performs significantly worse than both R_2C and the two-server competition approach of R_2G. Interestingly, Fig. 5 shows that the revenue rates for the R_2C and R_2G systems are nearly identical as functions of normalized arrival rate λ. We further investigate this observation using a conventional loss metric from game theory, the *price of anarchy*, which quantifies the performance loss due to decentralized competition relative to the centralized benchmark, defined as

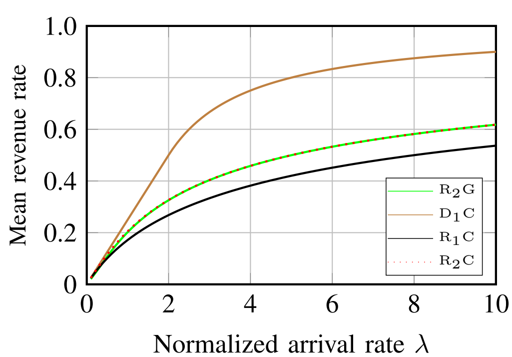

Fig. 5: Growth of mean revenue rate with arrivals.

$$\mathcal{A} \triangleq 1 - \frac{R_{R_2G}}{R_{R_2C}}.$$

As shown in Fig. 6, there is a measurable difference in revenue rate when normalized arrival rate λ is close to 0. However, this difference diminishes as the arrival rate λ approaches unity, and essentially disappears for $\lambda > 2$.

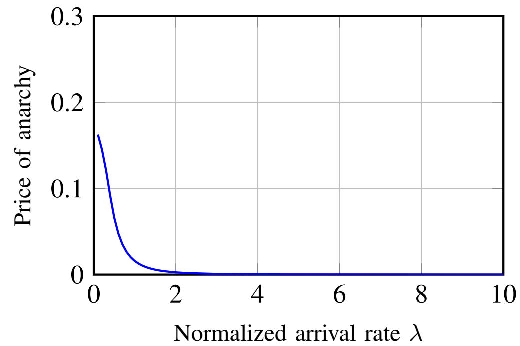

Fig. 6: Price of anarchy as a function of arrival rate.

Mean Price: Fig. 7 shows the price set by the server for the different systems under consideration. For normalized arrival rate $\lambda \geqslant 2$, the price under D_1C increases linearly in the arrival rate λ, while under R_1C, it grows proportionally to $\sqrt{\lambda}$. The figure suggests that the pricing behavior in R_2C and R_2G follow a trend slightly above $\sqrt{\lambda}$ as the normalized arrival rate λ increases.

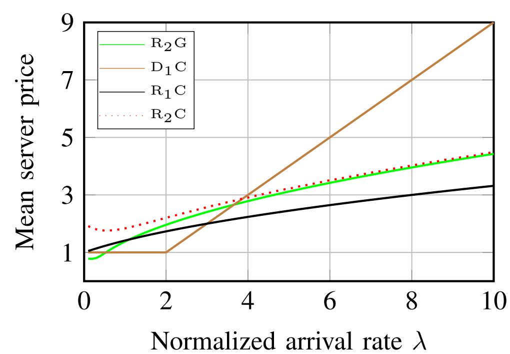

Fig. 7: Growth of mean server price with arrivals.

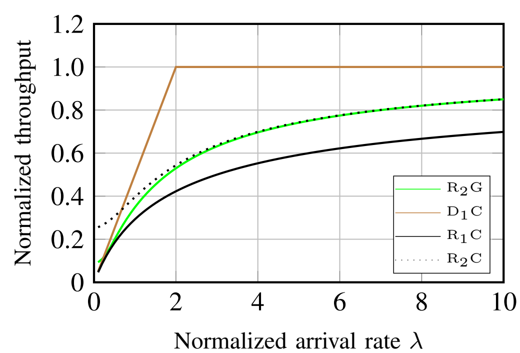

Fig. 8: Growth of throughput with arrivals.

Normalized Throughput: Fig. 8 presents the system throughput normalized by the number of servers. We observe that R_2G and R_2C closely track the performance of the fully centralized system D_1C for normalized arrival rates $\lambda < 1$. As the load increases, the blocking probability inherent in

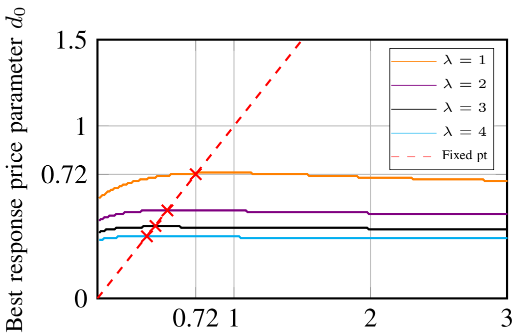

Fig. 9: Mean field game equilibrium.

randomized matching also rises, leading to a reduction in throughput, which is upper-bounded by unity in any case.

These results indicate that revenue performance of power-of-two randomized matching in the game-theoretic setting is very close to that of the one with collaborative pricing. This indicates that the loss in revenue is primarily due to the suboptimal routing in randomized matching systems when compared to centralized systems.

B. Price Determination

We briefly outline a set of numerical tools that network designers can use to implement our proposed pricing scheme. A numerical approach for each agent to compute a solution is via the following procedure.

1. **Compute z^*:** Given the normalized arrival rate λ and price parameter d_1 for all other agents, there exists a unique mean field limit of agent occupancy $z^*(d_1)$ as defined by (12).
2. **Compute best response $d_0(d_1)$:** Using the mean field limit $z^*(d_1)$, compute the corresponding self revenue rate via (17). Agent's best response is the price parameter $d_0(d_1)$ that maximizes the self revenue rate.
3. **Interpolation:** Evaluate $d_0(d_1)$ over a range of values for price parameter d_1 and construct an interpolated function. In our analysis, we use linear interpolation, although more advanced methods can be employed as needed.
4. **Find the fixed points:** A fixed point of best response satisfies $d_0 = d_1$. Thus, the mean field game equilibrium corresponds to the intersection of the curve $d_0(d_1)$ with the identity line $d_0 = d_1$.

An illustration of this process is shown in Fig. 9. We evaluated the mean-field game equilibrium for normalized arrival rates in the range $\lambda \in \{1, \dots, 4\}$. The corresponding mean-field game equilibrium values are $d_0^* = \{0.72, 0.51, 0.42, 0.36\}$, which are marked in Fig. 9.

C. Simulations for Different Noise Distributions

To evaluate our system under more general pricing models, it is necessary to use a simulator or oracle capable of computing the revenue rate and the mean field for a fixed price d_1 for all untagged servers. With such a simulator, the procedure described in the previous subsection can be applied to compute the mean-field game equilibrium.

We simulate a system with $N = 100$ servers to explore the effects of different price and valuation models. Simulations are conducted for three types of price and valuation distributions: Gaussian, Beta, and Log-normal. Servers draw the price and value from the identical family of distributions. When servers draw prices or valuations from parametric distributions with multiple parameters, we allow only one parameter (typically the mean) to vary, while fixing the others. For instance, we fix the variance at 1 in the Gaussian and Log-normal cases, and fix the shape parameter $b = 1$ in the Beta case. Servers adjust only the means based on observations. We simulate the evolution of the N server system under these general price and valuation distributions. For fixed valuation and price parameter for the untagged servers, we study the evolution of ensemble average of server occupancy and find the limiting time average of revenue rate for the tagged server.

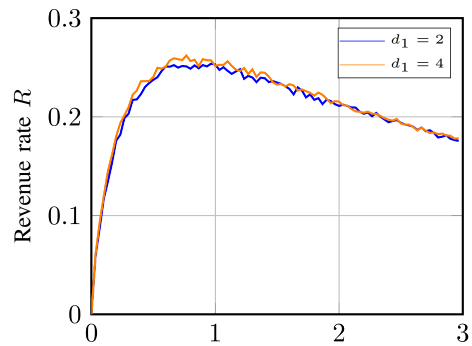

Fig. 10: Unimodality of the revenue rate at the tagged server.

An example under the Beta distribution is shown in Fig. 10 for unit mean value and normalized arrival rate of $\lambda = 5$. We observe that the revenue rate is unimodal for this distribution family. The same observation holds for other two distribution families as well, though we omit the figures for brevity. Table I reports the mean-field game equilibrium of occupancy and price pair (z^*, d^*), and the corresponding revenue rate $R(z^*)$ at the equilibrium, for different price and valuations distributions for unit mean value and normalized arrival rate $\lambda = 5$. These results illustrate that the mean-field game equilibrium can be explicitly computed using a simulator and the insights derived from analytically tractable cases continue to hold in the more realistic pricing environments.

Price/value distribution	z^*	d^*	$R(d^*)$
Exponential	0.690	0.31	0.4998
Gaussian	0.423	1	0.506
Beta	0.4256	1.034	0.2552
Log-normal	0.4282	0.933	2.233

TABLE I: Mean-field game equilibrium.

VII. Conclusion

In this paper, we examined the dynamics of a large-scale service marketplace and proposed a scalable, decentralized job matching and pricing mechanism based on randomized two-server selection and competitive pricing. Our model captures the essential trade-offs between centralized efficiency and decentralized implementability. We showed that simply increasing the choice set to two servers per job significantly improves performance in terms of blocking probability and system revenue, even without centralized coordination.

Our analysis established the existence and convergence of a mean field game equilibrium, ensuring stable and predictable system behavior in the large-system limit. The proposed $\mathrm{R_2G}$ mechanism was shown to deliver revenue scaling that closely matches centralized benchmarks such as $\mathrm{R_2C}$, while being simpler to implement and more robust.

Through extensive numerical and simulation results, we validated our analytical insights and further demonstrated that pricing insights derived from the exponential case remain effective for broader valuation models, including Gaussian, Beta, and Log-normal distributions. This confirms the practical applicability of our approach to more general environments.

Overall, our results suggest that a small amount of randomization in server choice, combined with strategic pricing via mean field game formulations, can form the backbone of efficient, decentralized cloud service marketplaces.

REFERENCES

[1] D. Narasimha, S. Nomula, S. Shakkottai, and P. Parag, "The power of two in large service marketplaces," in *IEEE Conf. Comp. Commun. (INFOCOM)*, May 2025, pp. 1–10.

[2] A. A. Cournot, "Researches into the mathematical principles of the theory of wealth," *Forerunners of Realizable Values Accounting in Financial Reporting*, 1838.

[3] I. Fisher, "Cournot and Mathematical Economics," *The Quarterly Journ. of Econ.*, vol. 12, no. 2, pp. 119–138, Jan. 1898.

[4] P. Naor, "The regulation of queue size by levying tolls," *Econometrica*, vol. 37, no. 1, pp. 15–24, Jan. 1969.

[5] N. M. Edelson and D. K. Hilderbrand, "Congestion tolls for poisson queuing processes," *Econometrica*, vol. 43, no. 1, pp. 81–92, Jan. 1975.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] ——, "Optimal pricing in a single server system," *ACM Trans. Model. Perform. Eval. Comput. Syst.*, vol. 8, no. 4, pp. 1–32, Dec. 2023.

[13] J. Kim and R. S. Randhawa, "The value of dynamic pricing in large queueing systems," *Op. Res.*, vol. 66, no. 2, Mar–Apr 2018.

[14] E. A. Feinberg and F. Yang, "Optimal pricing for a $GI/M/k/N$ queue with several customer types and holding costs," *Que. Sys. Theory Appl. (QUESTA)*, vol. 82, no. 1–2, pp. 103–120, Feb. 2016.

[15] 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.

[16] 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," *ACM Comput. Surveys*, vol. 55, no. 8, pp. 1–41, Aug. 2023.

[17] J.-M. Lasry and P.-L. Lions, "Mean field games," *Japanese J. Math.*, vol. 2, pp. 229–260, Mar. 2007.

[18] P. E. Caines, M. Huang, and R. P. Malhamé, "Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle," *Commun. Info. Sys.*, vol. 6, no. 3, pp. 221–252, Jan. 2006.

[19] 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.

[20] 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.

[21] 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.

[22] 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.

[23] ——, "Age-dependent distributed MAC for ultra-dense wireless networks," *IEEE/ACM Trans. Netw.*, vol. 31, no. 4, pp. 1674–1687, Aug. 2023.

[24] S. Aggarwal, M. A. uz Zaman, M. Bastopcu, S. Ulukus, and T. Başar, "Distributed offloading in multi-access edge computing systems: A mean-field perspective," *IEEE Trans. Mobile Comput.*, vol. 25, no. 6, pp. 7592–7610, Jun. 2026.

[25] 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, Nov. 2022, pp. 36 243–36 258.

[26] 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.

[27] B. Xia, S. Shakkottai, and V. Subramanian, "Small-scale markets for bilateral resource trading in the sharing economy," in *IEEE Conf. Comp. Commun. (INFOCOM)*, Apr. 2018, pp. 2447–2455.

[28] A. L. Stolyar, "Large-scale heterogeneous service systems with general packing constraints," *Adv. App. Prob.*, vol. 49, no. 1, pp. 61–83, Mar. 2017.

[29] A. Mukhopadhyay, A. Karthik, R. R. Mazumdar, and F. Guillemin, "Mean field and propagation of chaos in multi-class heterogeneous loss models," *Perf. Eval.*, vol. 91, no. C, pp. 117–131, Sep. 2015.

[30] Q. Xie, X. Dong, Y. Lu, and R. Srikant, "Power of d choices for large-scale bin packing: A loss model," *ACM SIGMETRICS Perf. Eval. Rev.*, vol. 43, no. 1, pp. 321–334, Jun. 2015.

[31] 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.

[32] 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, Jan. 1996.

[33] L. Ying, "On the approximation error of mean-field models," in *ACM SIGMETRICS Inter. Conf. Meas. Model. Comp. Sci.*, Jun. 2016, pp. 285–297.

[34] J.-Y. L. Boudec, D. McDonald, and J. Mundinger, "A generic mean field convergence result for systems of interacting objects," in *Intern. Conf. Quant. Eval. Syst. (QEST)*, vol. 4, no. 1, Sep. 2007, pp. 3–18.

[35] M. Benaïm and J.-Y. L. Boudec, "A class of mean field interaction models for computer and communication systems," *Perf. Eval.*, vol. 65, no. 11–12, pp. 823–838, Nov. 2008.

[36] R. W. Wolff, "Poisson arrivals see time averages," *Op. Res.*, vol. 30, no. 2, pp. 223–231, Mar-Apr 1982.

[37] C. Berge, *Topological Spaces: Including a treatment of multi-valued functions, vector spaces, and convexity*. Oliver and Boyd, 1963.

[38] L. E. J. Brouwer, "Über abbildung von mannigfaltigkeiten," *Mathematische Annalen*, vol. 71, pp. 97–115, Mar. 1911.

[39] J. Doncel, N. Gast, and B. Gaujal, "Are mean-field games the limits of finite stochastic games?" *ACM SIGMETRICS Perf. Eval. Rev.*, vol. 44, no. 2, pp. 18–20, Sep. 2016.

[40] A. K. K.S., C. K. Singh, S. T. Maguluri, and P. Parag, "Optimal pricing in multi server systems," in *Inter. Symp. on Model. Opt. Mobile, Ad Hoc, Wirel. Netw. (WiOPT)*, vol. 18, no. 1, Jun. 2020, pp. 1–8.

[41] 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, Dec. 1999.

[42] A. Coddington and N. Levinson, *Theory of Ordinary Differential Equations*, ser. International series in pure and applied mathematics. McGraw-Hill, 1955.

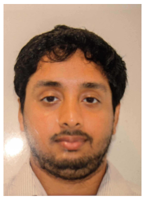

Dheeraj Narasimha obtained his PhD degree from Arizona State University working on mean-field games for density dependent continuous time Markov processes. He is currently a post-doctoral fellow at LAAS CNRS, Toulouse. His work on mean-field games has won runner-up best-paper prizes in ACM MobiHoc 2019 as well as most recently in IEEE INFOCOM 2025. The main focus of his work revolves around solving decision problems with many participants under weak interaction, be it through weak graph limits, potential games or mean field game structures.

Srinivas Nomula (S'22) is a PhD candidate in the Electrical Communication Engineering department at Indian Institute of Science, Bengaluru since Fall 2022. Prior to that, he received his B.Tech. degree in Electronics and Communication Engineering from the Indian Institute of Information Technology, Trichy in 2022. His areas of interest include stochastic modeling, distributed computing, sustainable computing, and hierarchical inference, in heterogeneous systems.

Srinivas Shakkottai (S'00-M'08-SM'15) received a PhD in 2007 in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign. He was a postdoctoral scholar in Management Science and Engineering at Stanford University in 2007. He joined Texas A&M University in 2008, where he is currently the Debbie and Dennis Segers '75 Professor at the Dept. of Electrical and Computer Engineering and the Dept. of Computer Science and Engineering (by courtesy). His research interests include wireless networks, reinforcement learning, caching and content distribution, multi-agent learning and game theory, networked markets, as well as data collection and analytics. He has received several national and international awards and served as area chair for AI/ML fora (NeurIPS, ICLR, AAAI) and networking fora (IEEE Infocom, ACM MobiHoc) and editor for IEEE Transactions on wireless communications as well as IEEE/ACM Transactions on networking.

Parimal Parag (S'04-M'11-SM'20) is an associate professor in the electrical communication engineering department at Indian Institute of Science, where he has been since Dec 2014. Prior to that, he was a senior system engineer (R&D) at Assia Inc. in Redwood City (2011-2014). He received Ph.D. degree from the Texas A&M University in 2011, the M.Tech. and B.Tech. degrees in 2004 from IIT Madras, all in electrical engineering. He is a coauthor of the 2018 IEEE ISIT student best paper, and a recipient of the 2017 early career award from the Science and Engineering Research Board. His research interests include the design and analysis of large scale distributed systems.

Appendix A
Proofs from Section III

A. Proof of Lemma 1

Let X and Y be two independent exponential random variables with rates a and b respectively, then we observe that $X \wedge Y$ is exponentially distributed with rate $a+b$, and

$$P\{X < Y\} = \mathbb{E}e^{-bX} = \frac{a}{a+b}. \tag{29}$$

Recall that valuation V_k and the tagged server price $P_{k,0}$ are independent exponential random variables with rates unity and d_0 respectively. Further untagged server prices $(P_{k,n} : n \in [N])$ are *i.i.d.* exponentially distributed with common rate d_1 and independent of V_k and $P_{k,0}$. Application of (29) to V_k and $P_{k,0}$ yields the result for q_1. Similarly, we obtain p_1 for V_k and $P_{k,n}$. Since $P_{k,n} \wedge V_k$ is exponentially distributed with rate $d_1 + 1$ and is independent of $P_{k,0}$, we obtain q_{20} from (29). Similarly, we can obtain q_{21} and p_2.

B. Proof of Lemma 2

Tagged server 0 is selected for service by the kth 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.

C. Proof of Lemma 3

Recall that $\sigma(X_{A_k,n}) \subseteq \mathcal{F}^N_{A_k}$ for all $n \in \mathcal{N}$ and I_k is selected independently and uniformly at random at each arrival instant. In particular, for each $n \in [N]$, we have

$$P\Big(\{I_k = \{0,n\}\} \mid \mathcal{F}^N_{A_k}\Big) = \frac{2}{N(N+1)}.$$

Hence, it follows from the the law of total probability, that

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

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

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

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 $Z^N_{A_k}$.

D. Proof of Proposition 1

The joint process $(X_{t,0}, Z^N_t, t \geqslant 0)$ is a continuous time Markov chain since the arrivals are Poisson, service times are *i.i.d.* exponentially distributed, and tasks are routed uniformly at random at each arrival instant.

(a) Recall that each busy server has an *i.i.d.* exponentially distributed task completion with unit rate. If $X_{t,0} = x$ and $Z_t = z$, then we observe that the tagged server can change from busy to idle with a unit rate, and z can decrease by $\frac{1}{N}$ with rate Nz. That is, we have

$$Q^N_{(1,z),(0,z)} = 1, \qquad Q^N_{(x,z),(x,z-\frac{1}{N})} = Nz.$$

We further recall that inter-arrival times to this system form an *i.i.d.* exponential sequence with rate $(N+1)\lambda$. An arrival leads to a transition for the tagged server to change from idle to busy if it is selected for arrival, and this transition rate is given by

$$Q^N_{(0,z),(1,z)} = (N+1)\lambda(r_1q_1 + r_0q_{20}).$$

Similarly, an arrival leads to increase in z by $\frac{1}{N}$ if a non-tagged server is selected for arrival, and this transition rate is given by

$$Q^N_{(x,z),(x,z+\frac{1}{N})} = (N+1)\lambda(xr_0p_1 + \bar{x}r_0q_{21} + s_0p_2 + s_1p_1).$$

Substituting r_j for $j \in \{0,1\}$ and s_ℓ for $\ell \in \{0,1\}$ from Lemma 3, we get the result.

(b) For each N, we observe that the state space $\{0,1\} \times \mathcal{Z}_N$ of the Markov process $(X_{t,0}, Z^N_t, t \geqslant 0)$ is finite. Hence, to show positive recurrence of this process, it suffices to show the irreducibility. The irreducibility follows from the fact that transition rate from each state (x,z) to its adjacent states $(x, z \pm \frac{1}{N})$ and $(1-x, z)$ are positive for $d_0, d_1 \in (0,\infty)$ and finite N.

E. Proof of Theorem 1

Recall from Lemma 5 that $h : [0,1] \to \mathbb{R}$ is defined as $h(z) = \lambda\bar{z}(2zp_1 + \bar{z}p_2) - z$ for each $z \in [0,1]$.

(a) Substituting $p_1 = \frac{d_1}{1+d_1}, p_2 = \frac{2d_1}{1+2d_1}$, and $q_{21} = \frac{d_1}{1+d_0+d_1}$ from Lemma 1 in (5), and taking absolute values on both sides, we obtain

$$|h(z) - h_N(x,z)| = \frac{\lambda\bar{z}2d_1}{N(1+2d_1)}\left|\frac{xd_1}{1+d_1} + \frac{\bar{x}(d_1-d_0)}{1+d_0+d_1}\right|.$$

The result follows from the triangular inequality, the fact that $|d_1 - d_0| \leqslant d_0 + d_1$, and that $\frac{x}{1+x} \leqslant 1$ for all $x \in \mathbb{R}_+$.

(b) Taking derivative of h with respect to z, and substituting the positive constants K_λ and L_λ from Definition 7, we obtain for any $z \in [0,1]$

$$h'(z) = -K_\lambda - L_\lambda z < 0. \tag{30}$$

From (30), we observe that $h'(z)$ is affine with negative slope $-L_\lambda$, and hence it is L_λ Lipschitz. Since h' is affine with a negative slope, it achieves the maximum and the minimum at $z = 0$ and $z = 1$, respectively.

(c) From the definition of h, we observe that $h(0) = \lambda p_2 > 0$ and $h(1) = -1$. From (30), we know that h is always decreasing, and hence it has a unique zero at $z^* \in [0,1]$. Any rest point $z \in \mathcal{S}$ of the deterministic mean-field equation defined in Definition 5 satisfies $h(z) = 0$. Substituting the form of h from Lemma 5, rearranging the equation, and substituting the positive constants K_λ, L_λ from Definition 7, it follows that the unique rest point z^* is the positive root of the following quadratic equation

$$z^2 + 2z\frac{K_\lambda}{L_\lambda} - 2\frac{(K_\lambda - 1)}{L_\lambda} - 1 = 0.$$

(d) Recall the definition of error function $\varepsilon : \mathbb{R}_+ \times [0,1] \to \mathbb{R}$ in (7) and Mckean-Vlasov equation (6) for autonomous non linear system $\Phi : \mathbb{R}_+ \times [0,1] \to [0,1]$. Since z^* is a constant, we obtain

$$\frac{\partial \varepsilon(t,z)}{\partial t} = \frac{\partial \Phi_t(z)}{\partial t} = h(\Phi_t(z)) = h(\varepsilon(t,z) + z^*).$$

Since $h(z^*) = 0$ for a rest point z^*, we can write the right hand side of the above equation as $h(\varepsilon + z^*) - h(z^*)$. For the specific form of h from (4) and definition of L_λ in (8), we can rewrite

$$\frac{\partial \varepsilon(t,z)}{\partial t} = \varepsilon(1-z^*)\frac{L_\lambda}{2} - \lambda\varepsilon(2p_1 z + \bar{z}p_2) - \varepsilon. \quad (31)$$

From the form of h from Lemma 5, positivity of normalized arrival rate λ, probability $p_2 \in (0,1)$ for $d_1 \in (0,\infty)$, and $z \in [0,1]$, we observe that

$$h(z) > z((1-z)\frac{L_\lambda}{2} - 1).$$

Since $h(z^*) = 0$ and $z^* > 0$ for $d_1 \in (0,\infty)$, we observe that $(1-z^*)\frac{L_\lambda}{2} < 1$. Therefore, we have for $\varepsilon \geqslant 0$

$$\begin{aligned} \frac{\partial \varepsilon(t,z)}{\partial t} &\leqslant -\lambda\varepsilon(2p_1 z + \bar{z}p_2), \varepsilon \geqslant 0, \\ \frac{\partial \varepsilon(t,z)}{\partial t} &> -\lambda\varepsilon(2p_1 z + \bar{z}p_2), \varepsilon < 0. \end{aligned} \quad (32)$$

If we can combine the results above to obtain that,

$$\frac{\partial\left|\varepsilon(t,z)\right|}{\partial t} \leqslant -\lambda\left|\varepsilon\right|(2p_1 \wedge p_2),$$

since $2p_1 > p_2$, the result follows. We note that above inequality is well defined for $|\varepsilon(t,z)| > 0$ since absolute value function is differentiable everywhere except at 0. In the following, we show that the inequality holds at $\varepsilon(t,z) = 0$, as well. Suppose there exists a $t_0 \in \mathbb{R}_+$ such that $\varepsilon(t_0,z) = 0$. If not, then there is nothing to show. We just need to show that $|\varepsilon(t,z)|$ is partially differentiable with respect to t at t_0. To this end, we need to show

$$\frac{\partial\left|\varepsilon(t,z)\right|}{\partial t} = \frac{|\varepsilon(t_0+h,z)| - |\varepsilon(t_0,z)|}{h} = \frac{|\varepsilon(t_0+h,z)|}{h}$$

is well defined and bounded. From the fundamental theorem of calculus, triangle inequality, and (32), we obtain

$$|\varepsilon(t_0+h,z)| \leqslant h\lambda(2p_1+p_2)\sup_{t\in[t_0,t_0+h]}|\varepsilon(t,z)|.$$

From (32), we observe that $\varepsilon(t,z)$ is continuous in $t \in [t_0, t_0+h]$. Since supremum of a continuous function in a compact interval is bounded, we observe that $\frac{|\varepsilon(t_0+h,z)|}{h}$ is bounded above.

(e) From the definition of generator matrix Q^N in Proposition 1(a), we obtain that $\sum_{y,w} Q^N_{(x,z),(y,w)}|w-z|^2$ equals

$$\frac{1}{N}(h(z)+2z) + \frac{1}{N}(h_N(x,z) - h(z)) \quad (33)$$

for any (x,z). Since $h(z)+2z$ is a continuous and concave function of z, it achieves a unique finite maximum in $z \in [0,1]$, which we denote as $C_\lambda \triangleq \sup_{z\in[0,1]}(h(z)+2z)$. Further, since $\sup_{x,z}|h(z) - h_N(z)| \leqslant \frac{\lambda}{N}$, we obtain

$$\sum_{y,w} Q^N_{(x,z),(y,w)}|w-z|^2 \leqslant \frac{C_\lambda}{N} + \frac{\lambda}{N^2}.$$

F. Perturbation theory

Definition 9. For the evolution of autonomous nonlinear system Φ governed by Mckean Vlasov equation (6) under map $h : [0,1] \to \mathbb{R}$, we define two maps $k, e : \mathbb{R}_+ \times [0,1]^2 \to \mathbb{R}$ for initial conitions $w, z \in [0,1]$ and time $t \in \mathbb{R}_+$ as

$$k(t,z,w) \triangleq \frac{\partial \Phi_t(z)}{\partial z}(w-z), \quad (34)$$

$$e(t,z,w) \triangleq \Phi_t(w) - \Phi_t(z) - k(t,z,w). \quad (35)$$

Lemma 6. *Consider the maps k, e from Definition 9 and $w, z \in [0,1]$. If the derivative $h'(z) = -K_\lambda - L_\lambda z$ for some positive constants K_λ, L_λ, then for all $t \in \mathbb{R}_+$*

$$|k(t,z,w)| \leqslant |w-z|\,e^{-K_\lambda t}, \quad (36)$$

$$|e(t,z,w)| \leqslant \frac{L_\lambda(w-z)^2}{K_\lambda + |w-z|\,L_\lambda} e^{-(K_\lambda - |w-z|L_\lambda)t}. \quad (37)$$

Proof: Since h' is Lipschitz, it follows that h is Lipschitz. From the definition of k in (34) we have used (6) to express the evolution of autonomous nonlinear system Φ. Therefore, by Leibniz rule, we can exchange the integral and the derivative to write

$$k(t,z,w) = (w-z) + \int_0^t h'(\Phi_s(z))k(s,z,w)ds.$$

Taking partial derivatives with time on both sides, we obtain

$$\frac{\partial k(t,z,w)}{\partial t} = h'(\Phi_t(z))k(t,z,w). \quad (38)$$

From the hypothesis on h, we have $-(K_\lambda + L_\lambda) \leqslant h'(x) \leqslant -K_\lambda$ for all $x \in [0,1]$. It follows that $h'(\Phi_t(z))k \leqslant -K_\lambda k \mathbb{1}_{\{k\geqslant 0\}}$ and $-h'(\Phi_t(z))k \leqslant (K_\lambda + L_\lambda)k\mathbb{1}_{\{k\geqslant 0\}}$. Combining the two results, we obtain

$$\frac{\partial\left|k(t,z,w)\right|}{\partial t} \leqslant -K_\lambda\left|k(t,z,w)\right|. \quad (39)$$

Since $k(0,z,w) = (w-z)$ and $|w-z| \leqslant 1$ for all $w, z \in [0,1]$, the result (36) for time decay of $|k(t,z,w)|$ follows.

From the definition of e in (35), the evolution of autonomous nonlinear system Φ in (6), we obtain

$$\frac{\partial e(t,z,w)}{\partial t} = h(\Phi_t(w)) - h(\Phi_t(z)) - \frac{\partial k(t,z,w)}{\partial t}. \quad (40)$$

From the mean value theorem applied to differentiable function h, we can find $\xi \in [0,1]$ to define $z_t^\xi \triangleq \Phi_t(z) + \xi(\Phi_t(w) - \Phi_t(z)) \in [0,1]$ such that

$$h(\Phi_t(w)) - h(\Phi_t(z)) = h'(z_t^\xi)(\Phi_t(w) - \Phi_t(z)). \quad (41)$$

From the form of h' in the hypothesis, we get $h'(z_t^\xi) - h'(\Phi_t(z)) = -\xi L_\lambda(\Phi_t(w) - \Phi_t(z))$. Combining this fact, $\xi \in [0,1]$, definition of e in (35), and (41), in (40), we obtain

$$\frac{\partial e}{\partial t} = -\xi L_\lambda k^2 + (h'(z_t^\xi) - \xi L_\lambda k)e \leqslant L_\lambda k^2 + (h'(z_t^\xi) + L_\lambda|k|)e.$$

Since $|k(t,z,w)| \leqslant |w-z|\,e^{-K_\lambda t}$ for all $t \in \mathbb{R}_+$ from (36), and $\sup_{z\in[0,1]} h'(z) \leqslant -K_\lambda$ from the hypothesis, we obtain

$$\frac{\partial e}{\partial t} \leqslant L_\lambda(w-z)^2 e^{-2K_\lambda t} - (K_\lambda - |w-z|\,L_\lambda)e.$$

Denoting constant $\alpha_\lambda \triangleq K_\lambda - |w-z| L_\lambda$, and rearranging the terms, we obtain

$$\frac{\partial e e^{\alpha_\lambda t}}{\partial t} = (\frac{\partial e}{\partial t} + \alpha_\lambda e)e^{\alpha_\lambda t} \leqslant L_\lambda (w-z)^2 e^{(\alpha_\lambda - 2K_\lambda)t}.$$

Integrating both sides in $[0,t]$, observing that $e(0,z,w)=0$ for any $z,w \in [0,1]$, and rearranging terms,we get

$$e(t,z,w) \leqslant \frac{L_\lambda (w-z)^2}{2K_\lambda - \alpha_\lambda}(e^{-\alpha_\lambda t} - e^{-2K_\lambda t}).$$

The result follows from ignoring the negative term on the right hand side of the above inequality. ■

Definition 10 (Poisson equation)**.** Consider the autonomous nonlinear system $\Phi : \mathbb{R}_+ \times [0,1] \to [0,1]$ following the Mckean Vlasov equation of Definition 5 under the map $h : [0,1] \to \mathbb{R}$. Let z^* be a rest point of h, then we say that a map $g : [0,1] \to \mathbb{R}$ satisfies the Poisson equation if

$$g'(z)h(z) = (z-z^*)^2. \tag{42}$$

Proposition 4. *The unique solution of Poisson equation is the map $g : [0,1] \to \mathbb{R}$ defined for each $z \in [0,1]$ as*

$$g(z) \triangleq - \int_{s\in\mathbb{R}_+} (\Phi_s(z) - z^*)^2 ds. \tag{43}$$

Proof: Taking $z \triangleq \Phi_t(z_0)$ for some $z_0 \in [0,1]$, observing that $\Phi_s \circ \Phi_t = \Phi_{s+t}$, and substituting $u \triangleq s+t$ in (43),

$$g(\Phi_t(z_0)) = -\int_{u\geqslant t} (\Phi_u(z_0) - z^*)^2 du.$$

Taking derivative with respect to t on both sides, we obtain the Poisson equation. Since the right hand side of (42) is Lipschitz and bounded, it follows by the Picard-Lindelof theorem [42] that the solution is unique. ■

Proposition 5. *Let $g : [0,1] \to \mathbb{R}$ be the solution of the Poisson equation* (42) *and $w,z \in [0,1]$. If $K_\lambda > |w-z| L_\lambda$, then*

$$-(g(w) - g(z) - g'(z)(w-z)) \leqslant M(w-z)^2, \tag{44}$$

where $M \triangleq \left(\frac{1}{2K_\lambda} + \frac{3L_\lambda}{K_\lambda^2 - (w-z)^2 L_\lambda^2}\right)$ is a positive constant.

Proof: Recall that $\Phi_s(w) - \Phi_s(z) = e(s,z,w) + k(s,z,w)$ from the definition of e and k in (35) and (34) respectively. It follows that $\Phi_s(w) - z^* = e(s,z,w) + k(s,z,w) + \varepsilon(s,z)$ from the definition of ε in (7). Substituting this in the definition of g in (43), we get

$$g(w) - g(z) = -\int_{s\in\mathbb{R}_+} ((e+k)^2 + 2\varepsilon(e+k))ds. \tag{45}$$

Since $\Phi_t(\cdot)$ is Lipschitz and has a Lipschitz derivative, applying Leibniz rule to interchange integrals and derivatives and by substituting the defintion of k from (34) and ε from (7), we get

$$g'(z)(w-z) = -2\int_{s\in\mathbb{R}_+} \varepsilon(s,z)k(s,z,w)ds. \tag{46}$$

Subtracting (45) from (46), we get

$$-\Big(g(w) - g(z) - g'(z)(w-z)\Big) = \int_{s\in\mathbb{R}_+} (e^2 + k^2 + 2e(k+\varepsilon))ds.$$

From the definitions of $e(s,z,w), k(s,z,w)$, and $\varepsilon(s,z)$, we get

$$|e + 2k + 2\varepsilon| \leqslant |\varepsilon(s,w)| + |\varepsilon(s,z)| + |k(s,z,w)|.$$

Since $|\varepsilon(s,z)| \leqslant 1$ and $\sup_{s\in\mathbb{R}_+} |k(s,z,w)| \leqslant |w-z| \leqslant 1$ for all $w,z \in [0,1]$, we get

$$|e + 2k + 2\varepsilon| \leqslant 3. \tag{47}$$

From (36), we obtain that

$$\int_{s\in\mathbb{R}_+} k(s,z,w)^2 ds \leqslant \frac{(w-z)^2}{2K_\lambda}.$$

From (47) and (37), we obtain

$$\int_{s\in\mathbb{R}_+} (e^2 + 2e(k + \varepsilon(s,z)))ds \leqslant \frac{3L_\lambda (w-z)^2}{K_\lambda^2 - (w-z)^2 L_\lambda^2}$$

Adding the last two inequalities, we get the result. ■

Lemma 7. *Consider a positive recurrent CTMC over state space $\mathcal{X}$ with generator matrix $Q \in \mathbb{R}^{\mathcal{X}\times\mathcal{X}}$ and invariant distribution $\pi \in \mathcal{M}(\mathcal{X})$ and a map $g \in \mathbb{R}^{\mathcal{X}}$. Then,*

$$\sum_{x\in\mathcal{X}} \pi_x \sum_{y\in\mathcal{X}} Q_{x,y}(g(y) - g(x)) = 0. \tag{48}$$

Proof: Since $g(y) = g(x)$ when $x = y$, we can restrict the summation in the left hand side of (48) to be over all $x,y \in \mathcal{X}$ such that $y \neq x$. From the property of generator matrix, we have $\sum_{y\neq x} Q_{x,y} = -Q_{x,x}$. Thus, we have

$$-\sum_{x\in\mathcal{X}} \pi_x g(x) \sum_{y\in\mathcal{X}:y\neq x} Q_{x,y} = -\sum_{x\in\mathcal{X}} g(x)\pi_x Q_{x,x}. \tag{49}$$

Since π is an invariant distribution, we have $\sum_{x\in\mathcal{X}:x\neq y} \pi_x Q_{x,y} = -\pi_y Q_{y,y}$. Therefore,

$$\sum_{y\in\mathcal{X}} g(y) \sum_{x\in\mathcal{X}:x\neq y} \pi_x Q_{x,y} = \sum_{y\in\mathcal{X}} g(y)\pi_y Q_{y,y}. \tag{50}$$

Adding (50) and (49), we get the result. ■

G. Proof of Theorem 2

Let $g : [0,1] \to \mathbb{R}$ be the solution of the Poisson equation (42) as defined in (43). We can add and subtract $h_N(x,z)$ defined in (3) from $h(z)$ in the right hand side of the Poisson equation, and subtract $\sum_{(y,w)} Q^N_{(x,z),(y,w)}(g(w) - g(z))$ on both sides, to obtain

$$\begin{aligned} &|z - z^*|^2 - \sum_{(y,w)} Q^N_{(x,z),(y,w)}(g(w) - g(z)) \quad (51)\\ &= g'(z)(h(z) - h_N(x,z)) \\ &- \sum_{(y,w)} Q^N_{(x,z),(y,w)}\Big(g(w) - g(z) - g'(z)(w-z)\Big). \end{aligned}$$

Recall from Proposition 1 that $(X_{t,0}, Z^N_t, t \geqslant 0)$ is a positive recurrent continuous time Markov chain for each $N \in [N]$, with generator matrix Q^N and invariant distribution π^N. Since the distribution of Z^N_∞ is π^N, we get

$$\mathbb{E}(Z^N_\infty - z^*)^2 = \sum_{(x,z)\in\{0,1\}\times\mathcal{Z}_N} \pi^N_{x,z}(z-z^*)^2.$$

Applying Lemma 7 to the generator matrix Q^N, the invariant distribution π^N, and the solution g of the Poisson equation (42), we get

$$\sum_{(x,z)} \pi^N_{x,z} \sum_{(y,w)} Q^N_{(x,z),(y,w)}(g(w) - g(z)) = 0.$$

Substituting $z \triangleq Z^N_\infty$ in (51), and taking expectation with respct to Z^N_∞ on both sides, we obtain that

$$\mathbb{E}(Z^N_\infty - z^*)^2 = \sum_{(x,z)} \pi^N_{x,z} \Bigg(g'(z)(h(z) - h_N(x,z)) - \sum_{(y,w)} Q^N_{(x,z),(y,w)} \Big(g(w) - g(z) - g'(z)(w-z) \Big) \Bigg).$$

From asymptotically accurate mean field model condition from Theorem 1(a), we have $\sup_{x,z} |h(z) - h_N(x,z)| \leqslant \frac{\lambda}{N}$. From (46), the fact that $\sup_{s\in\mathbb{R}_+, z\in[0,1]} |\varepsilon(s,z)| \leqslant 1$, and (36), we get

$$|g'(z)(w-z)| \leqslant 2\int_{s\in\mathbb{R}_+} |\varepsilon(s,z)|\,|k(s,z,w)|\,ds \leqslant 2\int_{s\in\mathbb{R}_+} |w-z|\,e^{-K_\lambda s} ds = \frac{2\,|w-z|}{K_\lambda}.$$

That is, $\sup_{z\in\mathcal{Z}_N} |g'(z)| \leqslant 2/K_\lambda$ for $w \neq z$. This together with the result $\sup_{x,z} |h(z) - h_N(x,z)| \leqslant \lambda/N$ from Theorem 1(a), we obtain

$$\sum_{(x,z)} \pi^N_{x,z} g'(z)(h(z) - h_N(x,z)) \leqslant \frac{2\lambda}{NK_\lambda}. \tag{52}$$

From the form of generator matrix Q^N in Proposition 1(a), we recall that $Q^N_{(x,z),(y,w)}(w-z)^2 > 0$ if and only if $|w-z| = \frac{1}{N}$. Therefore, we have $K_\lambda - (w-z)L_\lambda > 0$ for state pair $(x,z),(y,w)$ such that the transition rate is non zero for sufficiently large N. From triangle inequality, Proposition 5, Equation (33) in the proof of bounded mean transition rate condition of Theorem 1(e), we get

$$\sum_{(y,w)} Q^N_{(x,z),(y,w)}(g(w) - g(z) - g'(z)(w-z)) \leqslant \Big(\frac{1}{2K_\lambda} + \frac{3L_\lambda}{K_\lambda^2 - \frac{1}{N^2}L_\lambda^2}\Big)\Big(\frac{\lambda}{N^2} + \frac{C_\lambda}{N}). \tag{53}$$

From the convergence to mean field model in Theorem 1(a) and the definition of C_λ in Theorem 1(e) in (53) and adding (53), we obtain the result.

H. Proof of Theorem 3

From the positive recurrence of $(X_{t,0}, Z^N_t) \in \{0,1\} \times \mathcal{Z}_N$ in Proposition 1, we obtain that the process converges in distribution to $\pi^N \in \mathcal{M}(\{0,1\}\times\mathcal{Z}_N)$. We note that Theorem 2 allows us to conclude that Z^N_∞ converges to z^*. However, this does not automatically imply that the joint state $(X_{\infty,0}, Z^N_\infty)$ will converge in N. The first part of our proof will establish this joint convergence result. We proceed by first establishing that the limit is well-defined and then using a fast simulation result of [34] to show convergence. The second part of the proof then establishes the explicit form of the distribution of the joint state.

We begin by noting that the equilibrium fraction of occupied tagged servers Z^N_∞ converges to z^* in the mean square sense, and hence in distribution, i.e. the marginal $\lim_{N\to\infty} \pi^N_z = \mathbb{1}_{\{z=z^*\}}$ for all $z \in [0,1]$. Further note that since the joint state $(X_{t,0}, Z^N_t)$ lies in a compact set $\{0,1\} \times [0,1]$, the set of invariant measures $\pi^{(N)}$ over the joints states are compact as well by Prokhorov's theorem. We can therefore conclude that any limit point $\pi \in \mathcal{M}(\{0,1\}\times[0,1])$ is well defined.

Next note, by [35, Section 3.2], one can always choose a sampling process to replace the continuous time model with a discrete time equivalent. In our case, we will use uniform sampling (uniformization) in order to construct a discrete time equivalent model such that the number of jumps per time step is bounded independent of N.[2] It now follows from the fast simulation result, [34, Theorem 5.1] that the marginal measure converges to a fixed limit point given the mean field limit, i.e. $\lim_{N\to\infty} \pi^N_{x|z^*} = \pi_{x|z^*}$. Further, from Theorem 2, since Z^N_∞ converges to z^*, we now have

$$\pi^N_{x,z} = \pi^N_{x|z}\pi^N_z \xrightarrow{N\to\infty} \pi_{x|z^*}\pi_{z^*} = \pi_{x,z^*}.$$

This completes the first part of the theorem. We now proceed to explicitly quantify $\pi_{x|z^*}$. Next, invoking again the fast simulation result [34, Theorem 5.1] and noting our mean field is a point mass z^*, one can view the tagged server as an independent continuous time Markov process *given the mean field limit*. Hence, one has

$$\pi_{0|z^*} + \pi_{1|z^*} = 1,$$
$$\pi_{0|z^*} Q^N_{(0,z^*)(1,z^*)} = \pi_{1|z^*} Q^N_{(1,z^*)(0,z^*)}.$$

Where the second equation follows from the global balance equation of the tagged server *given the mean field*. Solving the two equations simultaneously and noting that π_z is a point mass at z^* gives us the required result.

Appendix B
Proofs from Section IV

A. Properties of revenue rate

Definition 11. We define a function $f : \mathbb{R}_+ \to \mathbb{R}_+$ for each $x \in \mathbb{R}_+$ in terms of fixed parameters λ, d_1, as

$$f(x) \triangleq \frac{z^*}{1+x} + \frac{1-z^*}{1+x+d_1}. \tag{54}$$

where z^* is a function of normalized arrival rate λ and untagged server's price rate d_1 defined in (12) as

$$z^*(d_1,\lambda) = -\frac{K_\lambda}{L_\lambda} + \sqrt{\frac{K_\lambda^2}{L_\lambda^2} + 1 + 2\frac{(K_\lambda - 1)}{L_\lambda}},$$

where $2d_1(K_\lambda - 1) = L_\lambda$ and $K_\lambda = 1 + \frac{2\lambda d_1}{(1+d_1)(1+2d_1)}$.

Lemma 8. *Let* $f^{(n)}$ *denote the* n*th derivative of* f *defined in* (54)*, then the following statements hold true for each* $x \in \mathbb{R}_+$ *and* $n \in \mathbb{N}$.

[2] This is precisely how we are able to simulate the joint system using uniformization.

(a) The nth derivative is positive for even n and negative for odd n, and equals

$$\frac{(-1)^n}{n!}f^{(n)}(x) = \frac{z^*}{(1+x)^{n+1}} + \frac{1-z^*}{(1+x+d_1)^{n+1}}. \quad (55)$$

(b) The following linear recursion holds,

$$nf^{(n-1)} + (x+d_1+1)f^{(n)} = (-1)^n n! \frac{z^* d_1}{(1+x)^{n+1}}. \quad (56)$$

(c) The following nonlinear recursion holds,

$$\frac{n}{(n!)^2}(f^{(n+1)}f^{(n-1)} - (f^{(n)})^2) \geqslant (f^{(n)})^2. \quad (57)$$

Proof: Consider the map f defined in (54).

(a) We can diretcly compute the nth derivative of f. From (55), it follows that $f^{(n)}$ is positive for even n and negative for odd n.

(b) From (55), we can write the linear recursion for the nth derivative of f in (56).

(c) Another recursion for nth derivative of f is

$$\frac{n}{(n!)^2}(f^{(n+1)}f^{(n-1)} - (f^{(n)})^2) = (f^{(n)})^2 + \frac{z^*(1-z^*)(n+1)d_1^2}{(1+x)^{n+2}(1+x+d_1)^{n+2}} \geqslant (f^{(n)})^2.$$

Therefore, the nonlinear recursive equation in (57) follows.

■

Definition 12. We define the set of extreme points for mean revenue rate as $\mathcal{D} \triangleq \{d_0 \in \mathbb{R}_+ : R'(d_0) = 0\}$. Let $a \triangleq \frac{1}{\sqrt{3}-1}$ and define four intervals that partition $\mathbb{R}_+$

$$J_1 \triangleq (0, a], \qquad J_2 \triangleq (a, d_1+1], \quad J_3 \triangleq (d_1+1, \infty), \quad J_4 \triangleq (2(d_1+1), \infty).$$

Lemma 9. *For fixed normalized arrival rate λ, and untagged servers' price rate d_1, the revenue rate $R : \mathbb{R}_+ \to \mathbb{R}_+$ defined in* (17)*, the following are true.*

(a) $R(0) = R(\infty) = 0$.

(b) $R'(0) > 0$,

(c) $R'(x) < 0$ for any $x \in J_3$,

(d) $\lim_{x\to\infty} R'(x) = 0$, and

(e) $R''(x) < 0$ for $x \in \mathcal{D} \cap J_1$.

Proof: We can rewrite the limiting mean revenue rate function R defined in (17), in terms of function f defined in (54), fixed parameters λ, d_1 and price rate of tagged server x, as

$$\Big(f(x) + \frac{1}{2\lambda x}\Big)R(x) = -f'(x). \quad (58)$$

We can write the first derivative of revenue rate as

$$(f + \frac{1}{2\lambda x})R' + (f' - \frac{1}{2\lambda x^2})R = -f''. \quad (59)$$

Multiplying both sides of (59) with $x^2(f + \frac{1}{2\lambda x})$, substiuting expression for f' in terms of R from (58), and rearranging terms, we obtain

$$\Big(xf + \frac{1}{2\lambda}\Big)^2 R' = x^2((f')^2 - f''f) - (xf'' + f')\frac{1}{2\lambda}. \quad (60)$$

(a) Multiplying both sides of (58) with x, we obtain $(xf(x) + \frac{1}{2\lambda})R(x) = -xf'(x)$. Setting $x = 0$, we obtain that $R(0) = 0$. From the form of f defined in (54), we obtain that $R(\infty) = 0$.

(b) It follows from (60) that $R'(0) = -2\lambda f'(0)$. Since $f'(0) < 0$ from (55), the result follows.

(c) From (57), we obtain that $f''f - (f')^2 \geqslant (f')^2$. Substituting this in (60), we get

$$(xf + \frac{1}{2\lambda})^2 R' \leqslant -(xf')^2 - (xf'' + f')\frac{1}{2\lambda}. \quad (61)$$

Substituting the first and the second derivative of f from (55) for $n = 1, 2$, we get

$$xf'' + f' = \frac{z^*(x-1)}{(1+x)^3} + \frac{(1-z^*)(x-1-d_1)}{(1+x+d_1)^3}.$$

It follows that $(xf'' + f') \geqslant 0$ for $x \in J_4$, and the result follows.

(d) From the expression for $(xf'' + f')$ in the previous part, we observe that $\lim_{x\to\infty}(xf'' + f') = 0$. Next, it follows from the expression for nth derivative of f in (55), we observe that $\lim_{x\to\infty} xf^{(n)}(x) = 0$ for all $n \in \mathbb{N}$. Further, we observe that $\lim_{x\to\infty} xf = 1$. Combining these three results and taking limit $x \to \infty$ on both sides of the upper bound on R' in (61), we get the result.

(e) Taking the derivative with respect to x on both sides of (60), we get that $(xf + \frac{1}{2\lambda})^2 R''(x)$ at any $x \in \mathcal{D}$ equals

$$2x((f')^2 - f''f) + x^2(f'f'' - f'''f) - (xf''' + 2f'')\frac{1}{2\lambda}.$$

Multiplying both sides of the above equation with x, using (60) for derivative of mean revenue rate, and the fact that $R'(x) = 0$ for $x \in \mathcal{D}$, cancelling terms, and rearranging them, we obtain that at any $x \in \mathcal{D}$

$$(f + \frac{1}{2x\lambda})^2 R''(x) = -(ff''' - f'f'') - \frac{(x^2 f''' - 2f')}{2\lambda x^3}. \quad (62)$$

To show the result, it suffices to show that $(ff''' - f'f'') > 0$ and $(x^2 f''' - 2f') > 0$ for all $x \in J_1$. From (55) for $n \in \{1, 2, 3\}$, we get that $f'''f - f'f''$ equals

$$\frac{4(z^*)^2}{(1+x)^5} + \frac{4(1-z^*)^2}{(1+x+d_1)^5} + \frac{4z^*(1-z^*)}{(1+x)^4(1+x+d_1)^4}\Big((1+x+d_1)^3 + (1+x)^3 + d_1^2(1+x+\frac{d_1}{2})\Big).$$

Since all the terms in the above equation is strictly positive for $x, d_1, \lambda > 0$, it follows that $f'''f - f'f'' > 0$. Next, we can write $-(x^2 f''' - 2f')$ in terms of $a = \frac{\sqrt{3}+1}{2}$ as

$$\frac{4z^*(x-a)(x+\frac{1}{2a})}{(1+x)^4} + \frac{4\bar{z^*}(x-(d_1+1)a)(x+\frac{(d_1+1)}{2a})}{(1+x+d_1)^4}.$$

It follows that $-(x^2 f''' - 2f') \leqslant 0$ for $x \leqslant a$, and hence the result follows.

■

B. Proof of Theorem 5

Recall from Lemma 9(a) that the limiting mean revenue rate $R(0) = R(\infty) = 0$. A sufficient condition to show that the mean revenue rate $R(d_0)$ is unimodal and has a maximum in an interval is to show that $R(d_0)$ is increasing and then decreasing in this interval. Since $R(d_0)$ is a differentiable function, it is equivalent to show that the derivative $R' \triangleq \frac{\partial R}{\partial d_0}$ has at most one downcrossing within the interval J_1.

First, we show that $R(d_0)$ has at least one local maxima in $J_1 = (0, d_1+1]$. This follows from the fact that mean revenue rate is increasing for zero tagged server price as shown in Lemma 9(b).

Next, we observe that the mean revenue rate is concave at all extreme points in J_1 as shown in Lemma 9(c). This implies that any extreme point of the mean revenue rate is its maxima, and hence a downcrossing. Therefore, over the interval J_1, either no crossing occurs or we have exactly one downcrossing. If no crossing occurs, then since $R'(0) > 0$, we have $R'(d_0) > 0$ for all $d_0 \in J_1$, and $d_0 = a$ is the unique revenue maximizing price over this interval. Otherwise, the maxima occurs at the unique downcrossing and the first part of the theorem is complete.

Finally, we observe that the mean revenue rate is decreasing for all points $x \in J_3 = (d_1+1, \infty)$ as shown in Lemma 9(e). $R''(d_0) > 0$ for any $d_0 \in \mathcal{D} \cap J_3$. We have completed the second part of the theorem.

Appendix C
Proofs from Section V

A. Proof of Proposition 3

Recall that servers have *i.i.d.* random price distributed exponentially with rate d and the arrivals have *i.i.d.* random valuations distributed exponentially with unit rate. Hence, the probability of an arrival finding a server's price below its valuation is $\frac{d}{d+1}$, and the thinned arrival rate to the system is $(N+1)\lambda\frac{d}{d+1}$. Suppose we set the common price parameter d such that the thinned arrival rates are restricted $\frac{d\lambda}{d+1} < 1$. For each time $t \in \mathbb{R}_+$, we let $X_{N+1}(t)$ denote the number of occupied servers, and we can define a normalized error term as

$$Y_{N+1}(t) \triangleq \frac{X_{N+1}(t) - (N+1)\frac{d\lambda}{d+1}}{\sqrt{N+1}}.$$

We can easily verify that X_{N+1} is a continuous time Markov chain, and hence so is normalized error Y_{N+1}. The Markov process Y_{N+1} converges to an Ornstein Uhlenbeck (OU) process [41] in distribution. The invariant measure of the limiting OU proces is a normal distribution with variance $\frac{d}{d+1}\lambda$. We can solve for $\frac{1}{N+1}X_{N+1}(t)$, to get for each time t

$$\frac{X_{N+1}(t)}{N+1} = \frac{Y_{N+1}(t)}{\sqrt{N+1}} + \lambda\frac{d}{d+1}.$$

Since Y_{N+1} is a zero mean Markov process with finite variance not scaling with N, it follows that $\frac{1}{N+1}X_{N+1}(t)$ converges in a mean square sense to a point mass at $\lambda\frac{d}{d+1}$ with rate $\frac{1}{\sqrt{N+1}}$. For *i.i.d.* random valuation distributed exponentially with unit rate and *i.i.d.* random price $P_{k,n}$ for arrival k at server n distributed exponentially with rate d, the system revenue rate is given by

$$(N+1)\lambda\mathbb{E}P_{k,n}\bar{G}(P_{k,n}) = (N+1)\lambda\frac{d}{(d+1)^2}.$$

It follows that normalized limiting per server revenue rate is $\lambda\frac{d}{(d+1)^2}$, and the maximum revenue rate for $\frac{d\lambda}{d+1} < 1$ is

$$R_{\mathrm{D_1C}} \triangleq \max\left\{\frac{\lambda d}{(d+1)^2} : \frac{d}{d+1} < \frac{1}{\lambda}\right\}.$$

Consider the unconstrained maximization of revenue rate $\frac{\lambda d}{(d+1)^2}$ as a function of price parameter d. We observe that this function is unimodal in d, with a unique maximum achieved at $d = 1$. If $\lambda < 2$, the maximizer $d = 1$ lies in the constraint set $\lambda\frac{d}{d+1} < 1$. Otherwise if $\lambda \geqslant 2$, then the revenue rate function is always increasing the price parameter d, and the maximum is achieved at the boundary of the constraint set for $d = \frac{1}{\lambda-1}$. Summarizing the two results we obtain that the revenue maximizing price rate

$$d_{\mathrm{D_1C}} = \begin{cases} 1, & \lambda < 2, \\ \frac{1}{\lambda-1}, & \lambda \geqslant 2. \end{cases}$$

Hence, we obtain the maximum revenue rate under $\mathrm{D_1C}$ as $R_{\mathrm{D_1C}} = \frac{\lambda d_{\mathrm{D_1C}}}{(d_{\mathrm{D_1C}}+1)^2}$. For the choice of $\lambda\frac{d}{d+1} < 1$, we have no busy servers in the limiting case of N being arbitrarily large. Hence, the blocking probability is the probability of price being higher than the valuation which is $\frac{1}{d_{\mathrm{D_1C}}+1}$. Hence the normalized throughput $\rho_{\mathrm{D_1C}} = \frac{\lambda d_{\mathrm{D_1C}}}{d_{\mathrm{D_1C}}+1}$.

B. Proof of Theorem 7

By symmetry, all servers have an identical evolution, and we can focus on server 0 without loss of generality. We can write the revenue rate at server 0 in terms of random price $P_{k,0}$ for kth arrival, as

$$\lim_{K\to\infty} \frac{1}{A_K}\sum_{k=1}^{K} P_{k,0}\xi_k,$$

where the selection indicator for server 0 can be written in terms of its server occupancy $X_{A_k,0}$ at the kth arrival time A_k and random valuation V_k for the kth arrival, as

$$\xi_k = \bar{X}_{A_k,0}\mathbb{1}_{\{V_k > P_{k,0}\}}.$$

Under random routing and *i.i.d.* exponential pricing in $\mathrm{R_1C}$ system, server 0 has a Poisson arrival rate of

$$\mathbb{E}\lambda\bar{G}(P_{k,0}) = \lambda\mathbb{E}e^{-P_{k,0}} = \frac{\lambda d}{d+1}. \tag{63}$$

Hence, the equilibrium probability of being idle for server 0 is denoted by $1 - z \triangleq \lim_{k\to\infty}\mathbb{E}\bar{X}_{A_k,0} = \frac{1}{1+\frac{\lambda d}{d+1}}$. Thus, the mean revenue rate is

$$\begin{aligned} R(d,\lambda) &= \lambda\lim_{k\to\infty}\mathbb{E}\bar{X}_{A_k,0}\mathbb{E}P_{k,0}\mathbb{1}_{\{V_k > P_{k,0}\}} \\ &= \frac{\lambda}{1+\frac{\lambda d}{d+1}}\mathbb{E}P_{k,0}\bar{G}(P_{k,0}) = \frac{\lambda d}{(d+1)^2(1+\frac{\lambda d}{d+1})}. \end{aligned}$$

In this case, the revenue is maximized for price parameter

$$d_{\mathrm{R_1C}} \triangleq \arg\max_{d} R(d,\lambda) = \frac{1}{\sqrt{1+\lambda}}.$$

The blocking probability due to server busyness for the price parameter $d_{\mathrm{R_1C}}$ is denoted by $z_{\mathrm{R_1C}}$. Conditioned on the server being idle, the blocking probability is due to valuation being lower than the server price, and is given by $\mathbb{E}G(P_{k,0})$.

$$p_b = z_{\mathrm{R_1C}} + (1 - z_{\mathrm{R_1C}})\mathbb{E}G(P_{k,0}).$$

Since idle probability is $(1 - z_{\mathrm{R_1C}}) = \frac{1}{1+\frac{\lambda d_{\mathrm{R_1C}}}{d_{\mathrm{R_1C}}+1}}$ and $\mathbb{E}\bar{G}(P_{k,0}) = \frac{d_{\mathrm{R_1C}}}{d_{\mathrm{R_1C}}+1}$ from (63), the throughput can now be written as follows

$$\rho_{\mathrm{R_1C}} = \lambda(1-p_b) = \lambda(1-z)\mathbb{E}\bar{G}(P_{k,0}) = \lambda\frac{d_{\mathrm{R_1C}}}{1+d_{\mathrm{R_1C}}(1+\lambda)}.$$