



Pigouvian Tolls and Welfare Optimality with Parallel Servers and Heterogeneous Customers

Tejas Bodas¹, Ayalvadi Ganesh^{2*}  and D. Manjunath³

Abstract | Congestion externalities are a well-known phenomenon in transportation and communication networks, healthcare etc. Optimization by self-interested agents in such settings typically results in equilibria which are sub-optimal for social welfare. Pigouvian taxes or tolls, which impose a user charge equal to the negative externality caused by the marginal user to other users, are a mechanism for combating this problem. In this paper, we study a non-atomic congestion game in which heterogeneous agents choose amongst a finite set of heterogeneous servers. The delay at a server is an increasing function of its load. Agents differ in their sensitivity to delay. We show that, while selfish optimisation by agents is sub-optimal for social welfare, imposing admission charges at the servers equal to the Pigouvian tax causes the user equilibrium to maximize social welfare. In addition, we characterize the structure of welfare optimal and of equilibrium allocations.

Keywords: Congestion games, Wardrop equilibrium, Discrete choice model, Pigouvian taxes

1 Introduction

We study service systems in which customers or agents can be served by any one of several heterogeneous servers. Customers arrive into the system according to a random process, reside in the system while being served, and then depart. Customers differ in their aversion to some congestion-based metric such as their sojourn-time in the system or the number of other customers with whom they share the server. We seek to determine how customers may be assigned to servers in such a way as to optimize some social welfare function, and also how pricing may be used to incentivize selfish customers to achieve the same social optimum.

Examples of such systems include web server farms, cloud and grid computing clusters, communication networks and cognitive radio systems. In these examples, customers may differ in the quality of service they require, and in their willingness to pay for it. The quality of service of a customer may depend on the share of bandwidth or other resources it receives, or the service latency or the sojourn time in the system. Another

example arises in transportation, where users may have a choice of tolled and toll-free routes, or between multiple modes of transport. Further examples include healthcare, where patients may be choosing between different service providers. Our modeling framework is quite general in this regard and encompasses all the above examples.

A common feature of the above examples is that the more customers choose a particular server, the worse their individual experience. For example, if more drivers choose a certain road, the slower the flow of traffic on it (above a certain utilization) and hence the longer the journey time. Similarly, if more patients choose a certain hospital, then they may have to wait longer for treatment, at least in the short run, when service capacities cannot be changed. This is known as a congestion externality.

Customer preferences are captured by a cost function that could depend on the system occupancy or sojourn time in a fairly arbitrary way. For example, in a transportation network, the cost function could be the expectation of a given function of the travel time, e.g., the probability

¹ Department of Electrical Engineering, IIT Dharwad, Dharwad, India.

² School of Mathematics, University of Bristol, Bristol, UK.

³ Department of Electrical Engineering, IIT Bombay, Mumbai, India.

*a.ganesh@bristol.ac.uk

that the travel time exceeds a certain threshold value. In a communication network, it could be a function of the bandwidth received, or the latency, or a combination of the two. We allow for customer heterogeneity by applying a suitable multiplier to the congestion cost. We call this multiplier its delay-sensitivity (but emphasise that congestion costs can take account of factors other than delay).

We do not constrain service policies except to insist that they be non-discriminatory and agnostic of customer characteristics. Thus, for instance, one server may adopt a first-come first-served (FCFS) policy while another splits its capacity equally amongst all its customers (processor-sharing or PS). Servers may charge a fixed admission price to each customer choosing that server; these can be different between servers but must be the same for each customer. In particular, servers cannot charge for priority or preferential treatment.

Customers choose a server so as to optimize their individual expected utility, i.e., to minimize the sum of the admission price and the expected congestion cost (weighted by their own delay-sensitivity). As the congestion cost depends on the choices of other customers, the interaction between them constitutes a game. The payoff structure makes this a congestion game^{21,27}. We assume in addition that customers are infinitesimal, i.e., that the impact of a marginal customer on the congestion cost at any server is negligible. This assumption renders the congestion game non-atomic. Nash equilibria in non-atomic congestion games are also known as Wardrop equilibria, from their origins in transportation networks³²; see (Nisan et al.²³, Chapter 18) for an overview of congestion games.

The goal of this paper is to study the social cost, i.e., the sum of congestion costs incurred by different customers weighted by their sensitivity to congestion, of a Wardrop equilibrium. In particular, we want to know if admission prices can be set in such a way as to ensure that the social cost at equilibrium is the minimum achievable by a central planner who could assign customers to servers. We answer this question in the affirmative. One set of such prices admit an interpretation as Pigouvian taxes associated with congestion externalities at the servers. While the welfare optimality of Pigouvian taxes is known in general, our contribution in this paper is to show that these depend only on the server, and not on the customer type. In other words, all customers using the same server are charged the same levy

(which may depend on the mix of customer types choosing that particular server).

A second contribution of the paper is a characterisation of the structure of socially optimal allocations and of Wardrop equilibria. Specifically, we show that in an optimal allocation, the server with the smallest congestion cost serves the most delay-sensitive customers, the one with the next smallest congestion cost serves the next most sensitive set of customers, and so on. We show that, for arbitrary admission prices at the servers, Wardrop equilibria have the same structure. Furthermore, the higher the admission price at a server, the lower its congestion cost (among servers that are utilized by some customer).

We survey some related work in the remainder of this section, before presenting a formal statement of the model and problem in the next, and stating our main results. Proofs are presented in the following section, and we conclude with a discussion of limitations of the current work and some open problems.

1.1 Related Work

The notion of a congestion externality was first formalized by Pigou²⁵, who proposed the use of a charge or levy to internalize the congestion externality in transportation networks, thereby guiding the system to a social optimum. Such charges are known as Pigouvian taxes and have since been studied in a wide variety of contexts including queueing systems^{8,18}, transportation networks^{30,33}, matching markets¹⁵ and climate change²⁰. While much of the work on Pigouvian taxes focuses on achieving socially optimal levels of consumption of a good associated with externalities, the work in this paper is most relevant when demand is inelastic (i.e., the quantity of demand does not depend on the price), but there is a choice between substitutes which generate different externalities. This is the case in many queueing and transportation applications. Secondly, our work considers heterogeneous agents, with different delay-sensitivities. In the following, we refer to them as multiclass customers, with “class” being used as a synonym for “delay-sensitivity”.

There is a substantial literature on the allocation of multi-class customers to parallel queues in both centralized and decentralized settings, including a variety of pricing schemes and game-theoretic formulations. Much of this work looks at specific cost functions arising from those models, whereas we consider a more general and abstract formulation. Below, we describe some

of the work more closely related to the approach taken in this paper and delineate these from the results we present. We use Kendall's notation for queueing models, which we now briefly describe. A queue is described by a triple $X/Y/n$, where X describes the arrival process, Y the job size distribution, and n the number of servers. Common choices for X are M , denoting Markovian and referring to a Poisson arrival process, and G , denoting "general" and referring to an arrival process in which the inter-arrival times are independent and identically distributed (i.i.d.), but with a general distribution. (Some authors prefer GI to emphasise the assumption of independence.) Common choices for Y are M , denoting Markovian and referring to job sizes with an exponential distribution, G , denoting i.i.d. job sizes with a general distribution, and D , denoting fixed, deterministic job sizes. If the service discipline is not the default $FCFS$ discipline, it is added to the notation. Thus, for example, an $M/G/1 - LCFS$ queue has Poisson arrivals, i.i.d. job sizes with a general distribution, and a single server which adopts a last-come-first-served policy.

There are several works that study the use admission prices to reduce congestion. Naor²², Edelson and Hilderbrand¹⁰ and Littlechild¹⁸ studied $M/M/1$ queues with identical customers who must choose between paying an admission price to enter the queue, incurring a random delay and receiving a fixed reward for service, or balking (i.e., leaving without being served). Admission prices are set by an operator who seeks to maximize revenue. If customers can observe the queue length on arrival and base their balking decision on it, then the revenue-maximizing admission price exceeds the one that maximizes social welfare²². However, if customers cannot observe the queue but must base their decision on only the known arrival and service rates, then these two admission prices coincide^{10,18}. In the latter setting, Littlechild¹⁸ obtained the admission fee as a Pigouvian tax and showed that this will induce a socially optimal arrival rate. Bradford⁸ extended the results to multiclass customers, each with their own delay cost function and reward for service, and obtains the Pigouvian admission charge for each class that achieves the socially optimal allocation. The admission charge is independent of the queue from which the customer receives service but depends on its class, which means that the system needs to elicit information of the customer class. In contrast, admission charges in our model are calculated for each queue but are agnostic of the customer class.

The equilibrium allocation of customers in multiqueue systems was studied by Bell and Stidham⁵, and Haviv and Roughgarden¹⁴. Both works focused on homogeneous customers, i.e., a single customer class. Bell and Stidham⁵ studied a set of parallel $M/G/1$ queues which differ in their holding cost per unit time and in their mean service time. They established structural properties of a socially optimal allocation as well as of Wardrop equilibria. Restricting their attention to parallel $M/M/1$ queues, Haviv and Roughgarden¹⁴ obtained an upper bound on the price of anarchy (PoA), defined as the ratio of the total cost at the Wardrop equilibrium to that at the social optimum. In comparison, we consider multiclass customer populations and general cost functions.

Borst⁷ studied the probabilistic allocation of multiclass traffic to parallel $M/G/1$ queues so as to minimize a specific social cost function, namely the total mean waiting cost per unit of time. He established a structural property of the optimal allocation. The structure we obtain for the optimal allocation is essentially the same, but our results apply to a very general class of queueing models and cost functions; we also do not restrict to finitely many customer classes. In addition, we consider a game-theoretic setting of selfish optimization and determine a pricing mechanism that will achieve social optimality with selfish optimization.

Sethuraman and Squillante²⁹ considered a variant of this problem where, in addition to optimal routing, servers decide the order in which customers in a queue are served, depending on their class, so as to optimise social welfare. An alternative approach is to allow customers to purchase priorities^{2,3,17,19,26}; a comprehensive survey of these and other similar models is presented by Hassin and Haviv¹³. Our work differs in that we do not allow servers to discriminate between customers, as a consequence of which they do not need to elicit information about customer class. This may be more realistic in certain applications.

A number of works have studied specific applications in which pricing is used to achieve service differentiation by incentivising end users to segregate themselves on the basis of their willingness to pay for higher quality or lower delay. In particular, there is a substantial body of work proposing charging for differentiated services (Diffserv) in the Internet, and studying the resulting user strategies and equilibria; see^{6,9,16,24}, for example. Additional examples include queues³¹ and transport networks³⁴. There has also been work on models in which prices are dynamically adapted in response to observed demands¹²; it

is shown that if prices adapt sufficiently slowly, then the system converges to a Nash equilibrium. Finally, while the work presented in this paper focuses on parallel queues, there has been considerable work on general networks; see Roughgarden²⁸ for a detailed discussion of selfish routing and the PoA, and Fleischer *et al.*¹¹ for the analysis of equilibria in a very general network model.

2 Model and Results

Consider a system with N parallel channels for service, which we refer to as servers or queues. Customers arrive into the system according to a marked Poisson process with intensity $\eta \times F$; here, η denotes the arrival rate, and F the distribution of the arriving customer's class or delay-sensitivity. The only assumption we make about the distribution F is that its support is bounded away from zero and infinity, i.e., that there are constants $\beta_{\min} > 0$ and $\beta_{\max} < \infty$ such that $F(x) = 0$ for all $x < \beta_{\min}$, and $F(\beta_{\max}) = 1$. Arriving customers must either select or be allocated to one of the queues upon arrival. We assume that the allocation has to be made with no knowledge of current or past queue occupancies, or past arrival times or routing decisions. Such an assumption may be less realistic for centralized allocation than when customers make individual decisions. Nevertheless, imposing this assumption uniformly permits clearer comparison of the two settings. The structure of Wardrop equilibria can be very different if queue occupancies are known, and requires a separate analysis, which is a topic for future research. In general, providing additional information can make the Wardrop equilibrium worse for all agents¹!

Under the assumption that queue occupancies are unknown, it is natural to restrict attention to Markovian policies, which route customers to queues according to some fixed probability vector that may depend on the customer's class, but not on history. (If queue occupancies are known, policies are Markovian with respect to a larger state space which includes that information.) We assume that customers of all classes have the same job size distributions, and that, once they join a queue, they are treated identically within it. Consequently, we assume that the congestion cost associated with a queue depends only on the aggregate arrival rate into that queue (and its service capacity and policies), but not on the composition of those arrivals. We make this precise below.

Let η denote the Borel measure on $[0, \beta_{\max}] \subset \mathbb{R}_+$ defined on intervals by

$$\eta((a, b]) = \eta(F(b) - F(a)). \quad (1)$$

In other words, the measure of an interval $(a, b]$ is defined as the total arrival rate of customers whose class lies in this interval. As usual, the measures of all Borel sets are determined by those of intervals. All measures in this paper are non-negative, finite Borel measures.

Now, Markovian routing corresponds to a decomposition of the measure η as

$$\eta = \lambda_1 + \dots + \lambda_N, \quad (2)$$

where λ_j is a measure on $[\beta_{\min}, \beta_{\max}]$ for each $j = 1, \dots, N$; arrivals into the j th queue of customers with classes in $(a, b]$ constitute a Poisson process of rate $\lambda_j((a, b])$. We denote the total arrival rate into the j th queue, and the mean delay-sensitivity of arrivals into this queue, by

$$\lambda_j = \lambda_j([\beta_{\min}, \beta_{\max}]) \text{ and } \bar{\lambda}_j = \int_{\beta_{\min}}^{\beta_{\max}} \beta d\lambda_j(\beta), \quad (3)$$

respectively.

Next, we associate with each queue j a cost function $D_j(\cdot)$ which specifies the congestion cost generated by a given aggregate arrival rate; thus, $D_j(\lambda)$ is the congestion cost incurred by each customer when the arrival rate into queue j is λ . The cost could be the mean sojourn time, or some higher moment of it, or the probability of the sojourn time exceeding a specified threshold. Our only assumption is that each function D_j be monotone increasing, continuous, and continuously differentiable in the interior of its domain (the set of arrival rates for which D_j is finite), with strictly positive derivative. In particular, we assume that the domain of each D_j is either \mathbb{R}_+ or an interval of the form $[0, a)$, and that in the latter case, $\lim_{x \uparrow a} D_j(x) = +\infty$.

The assumptions above are rather mild. We do not restrict the number of servers at a queue or the service discipline. Indeed, different queues may have different numbers of servers and employ different service disciplines. They can also be associated with different cost functions, for example the mean sojourn time at one queue and the second moment at another. The only requirement is that each queue treat all customers alike, irrespective of their class. In addition to traditional queueing models, our set-up also encompasses transportation models, where the mean journey time on a road may be some increasing function of the traffic intensity on it. The main

motivation for the assumption of Poisson arrivals is that it makes each D_j a function of a single real variable. It is not obvious how the monotonicity and differentiability assumptions would generalize if D_j were to be a function of the law of a stochastic process.

We are now ready to state the social welfare maximization problem. The objective is

$$\inf_{\lambda_1, \dots, \lambda_N} U(\lambda_1, \dots, \lambda_N) = \sum_{j=1}^N \bar{\lambda}_j D_j(\lambda_j),$$

subject to $\lambda_1 + \dots + \lambda_N = \eta$. (4)

Thus, the social cost is defined as the sum of the expected costs incurred by customers of different classes at different queues, weighted by the corresponding flow rates.

Our first result states that, if the social cost minimization problem is feasible, then it has a solution, i.e., the minimum is attained.

Lemma 1 *Let η be a finite measure with bounded support. Suppose that the cost functions D_j , $j = 1, \dots, N$, satisfy the assumptions stated above. If the optimization problem in (4) is feasible, i.e., there is some decomposition $(\lambda_1, \dots, \lambda_N)$ of η such that $D_j(\lambda_j)$ is finite for all $j = 1, \dots, N$, then (4) has a solution $(\lambda_1^*, \dots, \lambda_N^*)$.*

Next, we consider the formulation of a game between customers. Here, we allow the queues to charge admission prices, denoted by c_j at queue j . The goal of a class β customer entering the system is to choose a queue j so as to minimize $c_j + \beta D_j(\lambda_j)$ where λ_j is determined through the strategies of all customers. We assume that the arrival intensity measure η and the cost functions $D_j(\cdot)$, $j = 1, \dots, N$ are common knowledge. As we assumed that customers do not have access to current or past queue occupancies, or the history of arrival times or routing choices, they are necessarily restricted to choosing a server according to a fixed probability distribution, albeit one that may depend on their class. Thus, once again, the joint strategies may be represented by a decomposition of the measure η into measures $\lambda_1, \dots, \lambda_N$. We want to know when such a decomposition corresponds to a Wardrop equilibrium of the game.

The condition for a decomposition $(\lambda_1, \dots, \lambda_N)$ of η to be a Wardrop equilibrium is that

$$c_j + \beta D_j(\lambda_j) \leq c_k + \beta D_k(\lambda_k) \quad \forall j, k = 1, \dots, N, \text{ and } \beta \in \text{supp}(\lambda_j), \quad (5)$$

where $\text{supp}(\eta)$ denotes the support of the measure η , namely the smallest closed set F such that $\eta(F^c) = 0$. Here, F^c denotes the complement of F . The condition in (5) roughly says that, if a positive mass of customers of class β , or in an arbitrarily small neighbourhood of it, use queue j , then the expected cost of a class β customer in that queue must be no higher than its expected cost in any other queue.

The existence of a Wardrop equilibrium can be shown by looking at an auxiliary optimization problem, following Beckmann et al.⁴ in the single-class setting, and Yang and Huang³⁴ in the multiclass setting with a finite number of classes. Consider the optimization problem

$$\inf \hat{U}(\lambda_1, \dots, \lambda_N) = \sum_{j=1}^N \left(\int_0^{\lambda_j} D_j(x) dx + c_j \int_{\beta_{\min}}^{\beta_{\max}} \frac{1}{\alpha} d\lambda_j(\alpha) \right),$$

subject to $\lambda_1 + \dots + \lambda_N = \eta$. (6)

The existence of a solution follows by Lemma 1. It can easily be shown that any solution satisfies (5), which are essentially first-order conditions for optimality in the auxiliary problem. We include a formal statement and proof for completeness.

Lemma 2 *The infimum in the optimization problem (6) is attained. Moreover, any minimizer $(\lambda_1^W, \dots, \lambda_N^W)$ is a Wardrop equilibrium, i.e., it satisfies the condition in (5).*

A natural mechanism design^A question is whether we can set admission prices in such a way that selfish users reacting to these prices would assign themselves to queues in the proportions required for optimizing social welfare. Our main result affirms that this is indeed the case if admission prices are set equal to Pigouvian taxes corresponding to a welfare-optimal allocation.

Theorem 1 *Let $(\lambda_1^*, \dots, \lambda_N^*)$ be a solution of the social cost minimization problem, (4). Set the admission price c_j at queue j to be*

^A Mechanism design deals with the problem of achieving desired social choice objectives by designing the rules of the game such that the socially desirable outcome is a Nash equilibrium of the game; see²³ for further details.

$$c_j = \bar{\lambda}_j^* D'_j(\lambda_j^*), \tag{7}$$

where D'_j denotes the derivative of D_j .

Then, $(\lambda_1^*, \dots, \lambda_N^*)$ is a Wardrop equilibrium, i.e., it satisfies (5) with these admission prices.

Notice that c_j given in (7) is precisely the total negative externality imposed on existing customers at this queue by the admission of a marginal customer, and is hence the Pigouvian toll for this queue.

We now turn to the question of computing the optimal decomposition of a given measure η . If we can compute the optimal allocation, then we can also compute the corresponding Pigouvian taxes. Note that we start by assuming that the measure η is given. In practice, one of the major challenges of implementing Pigouvian taxes is eliciting utility functions; in our context, that corresponds to eliciting the true delay sensitivities β of different agents. Getting agents to truthfully reveal their preferences is a major challenge in mechanism design, and one which we do not address in this paper. Instead, we restrict ourselves to computing the optimal allocation given the true distribution of delay sensitivities.

The constraint on $(\lambda_1, \dots, \lambda_N)$ in the optimization problem (4) is linear, and so the set of measures satisfying the constraint is convex. If the cost function $\sum_{j=1}^N \bar{\lambda}_j D_j(\lambda_j)$ were a convex function of $(\lambda_1, \dots, \lambda_N)$, then the optimization problem would be convex, and could be solved using gradient descent methods. Unfortunately, this is not necessarily the case, as illustrated by the following counterexample.

Consider a system with two classes of customers and two $M/M/1$ queues. Class i customers arrive according to a Poisson process of rate η_i and have delay sensitivity β_i . Thus, the arrival intensity measure is $\eta = \eta_1 \delta_{\beta_1} + \eta_2 \delta_{\beta_2}$, where δ_x denotes the Dirac delta which puts unit mass at x . The job sizes for both classes are assumed to be i.i.d. exponential random variables with unit mean. Both servers have a unit service rate. We assume that $\eta_1 + \eta_2 < 1$, so that all allocations are feasible.

Recall that the mean delay in an $M/M/1$ queue with arrival rate λ and service rate 1 is $1/(1 - \lambda)$. Hence, the (class-weighted) congestion cost corresponding to a decomposition (λ_1, λ_2) of η is given by

$$U(\lambda_1, \lambda_2) = \frac{\bar{\lambda}_1}{1 - \lambda_1} + \frac{\bar{\lambda}_2}{1 - \lambda_2}.$$

The constraint that λ_1 and λ_2 are non-negative and decompose η is equivalent to the constraints that $\lambda_1 + \lambda_2 = \eta_1 + \eta_2$, $\bar{\lambda}_1 + \bar{\lambda}_2 = \beta_1 \eta_1 + \beta_2 \eta_2$, and that they are all non-negative. Thus, the welfare optimization problem (4) can be rewritten as

$$\begin{aligned} \inf U(\lambda_1, \bar{\lambda}_1, \lambda_2, \bar{\lambda}_2) &= \frac{\bar{\lambda}_1}{1 - \lambda_1} + \frac{\bar{\lambda}_2}{1 - \lambda_2}, \\ \text{subject to } \lambda_1 + \lambda_2 &= \eta_1 + \eta_2, \\ \bar{\lambda}_1 + \bar{\lambda}_2 &= \beta_1 \eta_1 + \beta_2 \eta_2, \\ \lambda_1, \bar{\lambda}_1, \lambda_2, \bar{\lambda}_2 &\geq 0. \end{aligned} \tag{8}$$

We now have the following negative result.

Lemma 3 *The optimization problem in (8) is not convex.*

In view of the above lemma, it is not obvious how to numerically compute socially optimal allocations in general. Nevertheless, we show below that both socially optimal allocations and Wardrop equilibria possess nice structural properties. These might suggest efficient algorithms for finding optima and equilibria in the model studied here.

Theorem 2 *Let $(\lambda_1^*, \dots, \lambda_N^*)$ achieve the minimum in (4). Suppose i and j are distinct queues, $\beta_2 > \beta_1 \geq 0$, and*

$$\lambda_j^*([\beta_2, \infty)) > 0 \text{ and } \lambda_i^*([0, \beta_1]) > 0.$$

Then $D_i(\lambda_i^) > D_j(\lambda_j^*)$. This inequality also holds if $\lambda_i^* = 0$ and $\lambda_j^* > 0$.*

The theorem says that if some of the customers served at queue j have higher delay sensitivity than some of the customers served at queue i (where “some” is to be interpreted as “a set of positive measure”), then the congestion cost at queue j must be smaller. Moreover, any queue which serves no customers (or a set of measure zero) must have larger congestion cost than any queue which serves some customers. The theorem implies that the queues segregate traffic by class as follows:

Corollary 1 *Suppose $(\lambda_1^*, \dots, \lambda_N^*)$ solves the optimization problem (4). Re-order the queues (permute their labels) such that $D_1(\lambda_1^*) \geq D_2(\lambda_2^*) \geq \dots \geq D_N(\lambda_N^*)$. Then, there exist $0 = \beta_0 \leq \beta_1 \leq \dots \leq \beta_N = \beta_{\max}$ such that $\text{supp}(\lambda_j^*) \subseteq [\beta_{j-1}, \beta_j]$ for all $j = 1, \dots, N$.*

The corollary says that customers are almost segregated by class, i.e., that each queue serves a

set of customer classes that is nearly disjoint from those served in other queues. By nearly disjoint, we mean that the customer classes served at distinct queues constitute intervals (closed, open or neither), which may only intersect at their boundaries. If the measure η has atoms (e.g., if there are only finitely many classes), then it is possible that customers belonging to some of these atoms are split across two or more queues. In routing terms, this would imply probabilistic routing to the corresponding queues. Secondly, the congestion costs at the queues are ordered such that more delay-sensitive customers incur smaller delays. Note that we are not claiming that queues with smaller delays have faster servers. Indeed, all servers may be identical, or the servers in less congested queues may even be slower! The differentiation in congestion costs is an emergent property of the optimal solution rather than a consequence of intrinsic differences between servers.

Next, we consider the same model, augmented with admission prices. Without loss of generality, we take $c_1 < c_2 < \dots < c_N$; if $c_i = c_j$, then we can collapse these two queues into a single queue whose delay function is the inf-convolution of the delay functions of its constituent queues, i.e.,

$$D(\lambda) = \inf\{D_i(\lambda_i) + D_j(\lambda_j) : \lambda_i, \lambda_j \geq 0, \lambda_i + \lambda_j = \lambda\}.$$

Each customer seeks to join a queue that minimizes the sum of the admission price, which is common to all classes, and the expected congestion cost, which is weighted by its own delay-sensitivity. We wrote down conditions in (5) for a decomposition of the arrival intensity measure η to be a Wardrop equilibrium. We now show that any Wardrop equilibrium has the same structure that we demonstrated above for a social optimum.

Theorem 3 *Suppose $(\lambda_1^W, \dots, \lambda_N^W)$ satisfies the conditions in (5), i.e., is a Wardrop equilibrium. Suppose i and j are distinct queues, $\beta_2 > \beta_1 \geq 0$, and*

$$\lambda_j^W([[\beta_2, \infty)) > 0 \text{ and } \lambda_i^W([[\beta_1, \infty)) > 0.$$

Then $c_j > c_i$.

The theorem says that if some of the customers served at queue j have higher delay sensitivity than some of the customers served at queue i , then the admission price at queue j must be larger. Whereas the social optimum does not use queues whose congestion cost at zero load is too

high, a queue could remain unused in a Wardrop equilibrium either because its congestion cost at zero load is too high, or because its admission price is too high, or a combination of the two. The theorem implies that the queues segregate traffic by class as follows:

Corollary 2 *Suppose $(\lambda_1^W, \dots, \lambda_N^W)$ satisfy the conditions in (5), with admission prices $c_1 < c_2 < \dots < c_N$. Then, there exist $0 = \beta_0 \leq \beta_1 \leq \dots \leq \beta_N = \beta_{\max}$ such that $\text{supp}(\lambda_j^W) \subseteq [\beta_{j-1}, \beta_j]$ for all $j = 1, \dots, N$.*

An important difference with the social optimum is that the ordering of queues by congestion cost at the social optimum is not obvious *a priori*. Hence, we do not know which queue will serve more delay-sensitive customers and which will serve less delay sensitive ones. On the other hand, at a Wardrop equilibrium, queues which charge a higher admission price (and are not idle) will serve more delay-sensitive customers than ones which charge a lower admission price.

3 Proofs

We now present proofs of the various results stated in the previous section.

Proof of Lemma 1 It is well-known that the set of sub-probability measures on \mathbb{R}_+ is compact in the weak topology. Hence, so too is the set of measures λ on \mathbb{R}_+ such that $\lambda \leq \eta$, where $\lambda = \lambda(\mathbb{R}_+)$, and $\eta = \eta(\mathbb{R}_+) < \infty$. By Tychonoff's theorem, the set $\{(\lambda_1, \dots, \lambda_N) : \lambda_i \leq \eta \forall i = 1, \dots, N\}$ is compact in the product topology. Next, the map $(\lambda_1, \dots, \lambda_N) \mapsto \lambda_1 + \dots + \lambda_N$ is continuous in this topology, and so the set $\{(\lambda_1, \dots, \lambda_N) : \lambda_1 + \dots + \lambda_N = \eta\}$ is closed. As it is a closed subset of a compact set, it is compact.

Let $\beta_{\max} = \sup\{\text{supp}(\eta)\}$. Then β_{\max} is finite by assumption. Hence, the support of λ_j is also restricted to $[0, \beta_{\max}]$ for all j , and the maps $\lambda_j \mapsto \bar{\lambda}_j$ are continuous in the weak topology; so, too, are the maps $\lambda_j \mapsto \lambda_j$, even without requiring bounded support. Finally, since the optimization problem (4) is feasible, we can restrict the minimization to a set of $(\lambda_1, \dots, \lambda_N)$ on which \mathcal{U} is bounded; in particular, each λ_j is in the domain of $D_j(\cdot)$. On this set, \mathcal{U} is continuous in the product topology. Thus, (4) involves the minimization of a continuous function over a compact set. Therefore, the minimum is attained. \square

Proof of Lemma 2 The constrained optimization problem (6) seeks the minimum of a continuous

function over a compact set; this follows along the same lines as the proof of Lemma 1. Hence, a minimizer exists.

Let $\lambda^W = (\lambda_1^W, \dots, \lambda_N^W)$ be one such minimizer. Suppose by way of contradiction that it is not a Wardrop equilibrium, i.e., that it does not satisfy (5). Then, there exist queues j and k such that

$$c_j + \beta D_j(\lambda_j^W) > c_k + \beta D_k(\lambda_k^W) \tag{9}$$

for some $\beta \in \text{supp}(\lambda_j^W)$.

By definition of the support, for any $\delta > 0$, there is an $\epsilon > 0$ such that $\lambda_j^W((\beta - \delta, \beta + \delta)) = \epsilon$. We now define a new decomposition of η which corresponds to shifting the mass in $(\beta - \delta, \beta + \delta)$ from queue j to queue k . More formally, denote the restriction of a measure μ to a set A by $\mu|_A$. Define $\mu = \lambda_j^W|_{(\beta-\delta, \beta+\delta)}$. For $\epsilon \in (0, 1)$, define

$$v_i^\epsilon = \begin{cases} \lambda_i^W, & i \neq j, k \\ \lambda_j^W - \epsilon \mu^{\beta, \delta}, & i = j, \\ \lambda_k^W + \epsilon \mu^{\beta, \delta}, & i = k. \end{cases}$$

Clearly, $v_i^\epsilon, i = 1, \dots, N$ are non-negative measures and decompose η , for any $\epsilon \in (0, 1)$. We see from (6) that

$$\begin{aligned} & \hat{U}(v^\epsilon) - \hat{U}(\lambda^W) \\ &= \int_{\lambda_k^W}^{\lambda_k^W + \epsilon \mu} D_k(x) dx - \int_{\lambda_j^W - \epsilon \mu}^{\lambda_j^W} D_j(x) dx \\ &+ \epsilon \int_{\beta - \delta}^{\beta + \delta} \frac{c_k - c_j}{\alpha} d\mu(\alpha) \\ &= \left(D_k(\lambda_k^W) - D_j(\lambda_j^W) + \frac{c_k - c_j}{\beta} \right) \mu\epsilon + o(\epsilon) + O(\delta\epsilon). \end{aligned}$$

By (9), the quantity in the last line above is negative, for small enough δ and ϵ . This contradicts the optimality of λ^W . The lemma is proved by contradiction. \square

Proof of Theorem 1 The proof is by contradiction. Suppose $\lambda^* = (\lambda_1^*, \dots, \lambda_N^*)$ solves the welfare optimization problem, (4), and that the admission prices c_j are set equal to the corresponding Pigouvian taxes, defined in (7). Suppose that $(\lambda_1^*, \dots, \lambda_N^*)$ do not satisfy (5), i.e., are not a Wardrop equilibrium for these prices. Then, there exist queues j and k such that

$$c_j + \beta D_j(\lambda_j^*) > c_k + \beta D_k(\lambda_k^*) \tag{10}$$

for some $\beta \in \text{supp}(\lambda_j^*)$.

By definition of the support, for any $\delta > 0$, there is an $\epsilon > 0$ such that $\lambda_j^*((\beta - \delta, \beta + \delta)) = \epsilon$. We now define a new decomposition of η which corresponds to shifting the mass in $(\beta - \delta, \beta + \delta)$ from queue j to queue k . Denoting the restriction of a measure μ to a set A by $\mu|_A$, we define

$$\lambda_i^{\beta, \delta} = \begin{cases} \lambda_i^*, & i \neq j, k \\ \lambda_j^* - \lambda_j^*|_{(\beta-\delta, \beta+\delta)}, & i = j, \\ \lambda_k^* + \lambda_j^*|_{(\beta-\delta, \beta+\delta)}, & i = k. \end{cases}$$

Clearly, $\lambda_i^{\beta, \delta}, i = 1, \dots, N$ are non-negative measures, and decompose η . We see from (4) that

$$\begin{aligned} & \mathcal{U}(\lambda^{\beta, \delta}) - \mathcal{U}(\lambda^*) \\ &= \bar{\lambda}_j^{\beta, \delta} D_j(\lambda_j^{\beta, \delta}) + \bar{\lambda}_k^{\beta, \delta} D_k(\lambda_k^{\beta, \delta}) \\ &\quad - \bar{\lambda}_j^* D_j(\lambda_j^*) - \bar{\lambda}_k^* D_k(\lambda_k^*) \\ &= \left(\bar{\lambda}_j^* - \beta\epsilon + O(\delta\epsilon) \right) \left(D_j(\lambda_j^*) - \epsilon D_j'(\lambda_j^*) + o(\epsilon) \right) \\ &\quad - \bar{\lambda}_j^* D_j(\lambda_j^*) \\ &\quad + \left(\bar{\lambda}_k^* + \beta\epsilon + O(\delta\epsilon) \right) \left(D_k(\lambda_k^*) + \epsilon D_k'(\lambda_k^*) + o(\epsilon) \right) \\ &\quad - \bar{\lambda}_k^* D_k(\lambda_k^*) \\ &= \epsilon \left(\beta D_k(\lambda_k^*) + \bar{\lambda}_k^* D_k'(\lambda_k^*) - \beta D_j(\lambda_j^*) - \bar{\lambda}_j^* D_j'(\lambda_j^*) \right) \\ &\quad + O(\delta\epsilon) + o(\epsilon). \end{aligned}$$

Substituting the expression for the Pigouvian taxes c_j and c_k from (7) in the above, we get

$$\mathcal{U}(\lambda^{\beta, \delta}) - \mathcal{U}(\lambda^*) = \epsilon (c_k + \beta D_k(\lambda_k^*) - c_j - \beta D_j(\lambda_j^*)) + O(\delta\epsilon) + o(\epsilon).$$

If we let δ decrease to zero, then so does ϵ , and the last two terms in the expression above are negligible compared to the first. Hence, it follows from the above and (10) that $\mathcal{U}(\lambda_1^{\beta, \delta}, \dots, \lambda_N^{\beta, \delta}) - \mathcal{U}(\lambda_1^*, \dots, \lambda_N^*) < 0$ for δ sufficiently small. This contradicts the assumed optimality of $(\lambda_1^*, \dots, \lambda_N^*)$.

We have thus shown by contradiction that the conditions, (5), for a Wardrop equilibrium must be satisfied at a socially optimal allocation when

the admission prices are given by Pigouvian taxes. \square

Proof of Lemma 3 The proof is an exercise in calculus. The set of $(\lambda_1, \bar{\lambda}_1, \lambda_2, \bar{\lambda}_2)$ satisfying the constraints in (8) is convex. A necessary condition for the objective function to be convex on the feasible set is that the Hessian of U be positive semi-definite on the subspace $\{(x_1, x_2, x_3, x_4) : x_1 + x_3 = 0, x_2 + x_4 = 0\}$ of feasible deviations, at each feasible point $(\lambda_1, \bar{\lambda}_1, \lambda_2, \bar{\lambda}_2)$.

Denoting the Hessian by $[D^2U]$, we consider the quadratic form

$$\begin{aligned} & (x_1, x_2, x_3, x_4)[D^2U(\lambda_1, \bar{\lambda}_1, \lambda_2, \bar{\lambda}_2)](x_1, x_2, x_3, x_4)^T \\ &= \frac{2\bar{\lambda}_1 x_1^2}{(1-\bar{\lambda}_1)^3} + \frac{2x_1 x_2}{(1-\bar{\lambda}_1)^2} + \frac{2\bar{\lambda}_2 x_3^2}{(1-\bar{\lambda}_2)^3} + \frac{2x_3 x_4}{(1-\bar{\lambda}_2)^2} \\ &= \left(\frac{2\bar{\lambda}_1}{(1-\bar{\lambda}_1)^3} + \frac{2\bar{\lambda}_2}{(1-\bar{\lambda}_2)^3} \right) x_1^2 \\ & \quad + \left(\frac{2}{(1-\bar{\lambda}_1)^2} + \frac{2}{(1-\bar{\lambda}_2)^2} \right) x_1 x_2, \end{aligned}$$

where we have used the fact that $x_1 = -x_3$ and $x_2 = -x_4$ on the subspace of interest to obtain the second equality. Now, it is clear that the expression above can be made negative by choosing x_1 and x_2 non-zero and of opposite signs, and x_1 sufficiently small in absolute value.

In other words, the quadratic form is not always non-negative, i.e., the Hessian is not positive semi-definite on the subspace of interest. Therefore, the objective function U is not convex on the feasible set. \square

Proof of Theorem 2 Let $\lambda^* = (\lambda_1^*, \dots, \lambda_N^*)$ solve (4), and let i, j, β_1 and β_2 be as in the statement of the theorem. We shall prove the theorem by contradiction.

Suppose first that $\lambda_i^* > 0$ and that $D_i(\lambda_i^*) < D_j(\lambda_j^*)$. We shall show that shifting a small mass of customer from queue j to queue i and an equal mass from i to j reduces the social cost, contradicting the optimality of λ^* . Let μ_i and μ_j be measures such that

$$\begin{aligned} & \mu_i \leq \lambda_i, \mu_j \leq \lambda_j, \mu_i = \mu_j > 0, \\ & \text{supp}(\mu_i) \subseteq [0, \beta_1], \text{supp}(\mu_j) \subseteq [\beta_1, \infty). \end{aligned}$$

It is clear from the assumptions that such measures exist. Since $\beta_1 > \beta_2$, we also have $\bar{\mu}_j > \bar{\mu}_i$.

Consider the measures $\tilde{\lambda}$ defined as follows:

$$\tilde{\lambda}_k = \begin{cases} \lambda_k^*, & k \neq i, j, \\ \lambda_i^* + \mu_j - \mu_i, & k = i, \\ \lambda_j^* - \mu_j + \mu_i, & k = j. \end{cases}$$

Then, $\tilde{\lambda}_k = \lambda_k^*$ for all k , since equal masses are swapped between queues i and j while flows into all other queues are unchanged. Hence, the congestion costs D_k at all queues remain unchanged. Thus, we get

$$U(\tilde{\lambda}) - U(\lambda^*) = (\bar{\mu}_j - \bar{\mu}_i)(D_i(\lambda_i^*) - D_j(\lambda_j^*)) < 0,$$

since $\bar{\mu}_j > \bar{\mu}_i$ as noted, while $D_i(\lambda_i^*) < D_j(\lambda_j^*)$ by assumption. But this contradicts the optimality of λ^* . Thus, we cannot have $D_i(\lambda_i^*) < D_j(\lambda_j^*)$ and $\lambda_i^* > 0$.

Suppose next that $\lambda_i^* > 0$ and $D_i(\lambda_i^*) = D_j(\lambda_j^*)$. Let $\tilde{\lambda}$ be as above, and define

$$\lambda^\alpha = \alpha \tilde{\lambda} + (1 - \alpha) \lambda^*, \quad \alpha \in [0, 1].$$

Then, for all $\alpha \in [0, 1]$, $\lambda_i^\alpha = \lambda_i^*$ and $\lambda_j^\alpha = \lambda_j^*$, so $D_i(\lambda_i^\alpha) = D_i(\lambda_i^*) = D_j(\lambda_j^*) = D_j(\lambda_j^\alpha)$. Hence, $U(\lambda^\alpha) = U(\lambda^*)$, which implies that $(\lambda_1^\alpha, \dots, \lambda_N^\alpha)$ solve the welfare optimization problem, (4), for every $\alpha \in [0, 1]$.

Now, for $\alpha \in (0, 1)$, and small enough $|\epsilon|$, define the measures $\mathbf{v}_k^{\alpha, \epsilon}, k = 1, \dots, N$, by

$$\mathbf{v}_k^{\alpha, \epsilon} = \begin{cases} \lambda_k^\alpha, & k \neq i, j, \\ \lambda_i^\alpha + \epsilon \mu_j, & k = i, \\ \lambda_j^\alpha - \epsilon \mu_j, & k = j. \end{cases}$$

If $|\epsilon|$ is sufficiently small, depending on α , then these are non-negative measures. We now have

$$\begin{aligned} & U(\mathbf{v}^{\alpha, \epsilon}) - U(\lambda^\alpha) \\ &= \bar{v}_i^{\alpha, \epsilon} D_i(v_i^{\alpha, \epsilon}) + \bar{v}_j^{\alpha, \epsilon} D_j(v_j^{\alpha, \epsilon}) \\ & \quad - \bar{\lambda}_i^\alpha D_i(\lambda_i^\alpha) - \bar{\lambda}_j^\alpha D_j(\lambda_j^\alpha) \\ &= \epsilon \left(\bar{\mu}_j D_i(\lambda_i^\alpha) + \mu_j \bar{\lambda}_i^\alpha D_i'(\lambda_i^\alpha) \right. \\ & \quad \left. - \bar{\mu}_j D_j(\lambda_j^\alpha) - \mu_j \bar{\lambda}_j^\alpha D_j'(\lambda_j^\alpha) \right) + o(\epsilon). \end{aligned}$$

For $U(\lambda^\alpha)$ to be a global minimum, the coefficient of ϵ in the above expression must be zero. Thus,

$$\mu_j \left(\bar{\lambda}_i^\alpha D_i'(\lambda_i^\alpha) - \bar{\lambda}_j^\alpha D_j'(\lambda_j^\alpha) \right) = \bar{\mu}_j (D_j(\lambda_j^\alpha) - D_i(\lambda_i^\alpha)).$$

But $\lambda_k^\alpha = \lambda_k^*$ for all $\alpha \in [0, 1]$ and $k = 1, \dots, N$. Combining this with the fact that $D_i(\lambda_i^*) = D_j(\lambda_j^*)$ by assumption, we can rewrite the last equation as

$$\mu_j \left(\bar{\lambda}_i^\alpha D_i'(\lambda_i^*) - \bar{\lambda}_j^\alpha D_j'(\lambda_j^*) \right) = 0. \tag{11}$$

Now, $\bar{\lambda}_i^\alpha$ is strictly increasing in α and $\bar{\lambda}_j^\alpha$ is strictly decreasing, as λ^α is obtained by swapping a

volume of more delay-sensitive traffic in queue j for an equal volume of less delay-sensitive traffic in queue i , and these volumes are increasing in α . Moreover, $D'_i(\lambda_i^*)$ and $D'_j(\lambda_j^*)$ are strictly positive, and hence non-zero, by assumption. It follows that (11) cannot hold for all $\alpha \in (0, 1)$, or even for two distinct values of α .

Thus, we have shown by contradiction that we cannot have $\lambda_i^* > 0$ and $D_i(\lambda_i^*) = D_j(\lambda_j^*)$. It only remains to consider the possibility that $\lambda_i^* = 0$. Let μ_j be as above. Fix $\epsilon > 0$ sufficiently small, and define the measures ν^ϵ as follows:

$$\nu_k^\epsilon = \begin{cases} \lambda_k^*, & k \neq i, j, \\ \epsilon \mu_j, & k = i, \\ \lambda_j^* - \epsilon \mu_j, & k = j. \end{cases}$$

Then, we have

$$\begin{aligned} \mathcal{U}(\nu) - \mathcal{U}(\lambda^*) &= \epsilon \mu_j (D_i(0) - D_j(\lambda_j^*)) \\ &\quad - \epsilon \mu_j \lambda_j^* D'_j(\lambda_j^*) + o(\epsilon). \end{aligned}$$

Since $D'_j(\lambda_j^*) > 0$, the above quantity is negative, contradicting the optimality of λ^* , unless $D_i(0) > D_j(\lambda_j^*)$. This completes the proof of the theorem. \square

Proof of Corollary 1 Suppose the corollary is false. Then, there is a solution $\lambda^* = (\lambda_1^*, \dots, \lambda_N^*)$ of (4), and queues i and j , such that $D_i(\lambda_i^*) \geq D_j(\lambda_j^*)$ but queue i also serves a non-zero mass of customers who are more delay-sensitive than some of the customers served in queue j . More precisely, there exist $\beta_2 > \beta_1$ such that $\lambda_i^*([\beta_2, \infty)) > 0$ and $\lambda_j^*([0, \beta_1]) > 0$. But this contradicts Theorem 2. \square

Proof of Theorem 3 Suppose $\lambda^W = (\lambda_1^W, \dots, \lambda_N^W)$ satisfies the conditions in (5). Suppose i and j are distinct queues and $\beta_2 > \beta_1 \geq 0$ are such that

$$\lambda_j^W([\beta_2, \infty)) > 0 \text{ and } \lambda_i^W([0, \beta_1]) > 0.$$

Pick $\beta \leq \beta_1 \in \text{supp}(\lambda_i^W)$ and $\gamma \geq \beta_2 \in \text{supp}(\lambda_j^W)$. We have by (5) that

$$\begin{aligned} c_i + \beta D_i(\lambda_i^W) &\leq c_j + \beta D_j(\lambda_j^W), \\ c_j + \gamma D_j(\lambda_j^W) &\leq c_i + \gamma D_i(\lambda_i^W). \end{aligned} \tag{12}$$

It follows from these inequalities that $(\gamma - \beta)(D_i(\lambda_i^W) - D_j(\lambda_j^W)) \leq 0$. Since $\gamma > \beta$, it follows that $D_i(\lambda_i^W) \geq D_j(\lambda_j^W)$. Substituting this in (12), we obtain that $c_i \leq c_j$. As it was assumed that admission prices are all distinct, we have $c_i > c_j$, as claimed. \square

Proof of Corollary 2 Consider two queues i and j . Suppose $\beta_1 \in \text{supp}(\lambda_i^W)$, $\beta_2 \in \text{supp}(\lambda_j^W)$ and $\beta_1 < \beta_2$. Then, there is a $\delta > 0$ sufficiently small that

$$\begin{aligned} \lambda_j^W([\beta_2 - \delta, \infty)) &> 0, \quad \lambda_i^W([0, \beta_1 + \delta]) > 0, \\ \beta - \delta &> \beta_1 + \delta. \end{aligned}$$

Hence, by Theorem 3, $c_j > c_i$, i.e., $j > i$. This proves the corollary. \square

4 Summary and Discussion

We considered a very general model of multiple parallel queues serving a heterogeneous customer population, and studied the problem of routing customers to queues so as to maximize social welfare. We characterized certain structural properties of the welfare-optimizing allocation. We also considered selfish routing decisions made by individual customers when the queues charge admission prices, and characterized the structure of Wardrop equilibria. Finally, we showed that, if the admission prices at the queues are set equal to the congestion externalities at a socially optimal allocation, then the social optimum coincides with a Wardrop equilibrium.

The setting we studied was very general, and encompassed a variety of applications with congestion externalities. Nevertheless, some of the assumptions are restrictive. We model customer heterogeneity by applying different multipliers to a common measure of congestion cost at each queue. But it might be the case that some customers care about mean delay, while others care about the probability of exceeding a certain threshold. In that case, no multiplier on the congestion cost would be appropriate for capturing this diversity. Another restrictive assumption is that customers may differ in delay sensitivity, but not in the distribution of the workload they bring into the system. Indeed, this is why Pigouvian tolls depend on the queue, but not on the customer class. If this assumption were relaxed, the externality imposed by a customer would depend on its workload, and hence on its class; this would need to be taken into account in setting Pigouvian tolls.

We briefly discussed the difficulty of determining the optimal allocation. We showed that the optimization problem is non-convex, but did not prove that it is hard. The structural properties of the optimal allocation that we established do not resolve this question, as the optimal ordering of the queues is unknown. Even if the optimal ordering were given, it is not entirely obvious

that the thresholds can be computed efficiently. Likewise, the computational complexity of determining the Wardrop equilibria is also unknown. Note that the ordering of queues in this case is determined by the given prices. Thus, one open problem for future research is developing efficient algorithms for these problems, or proving that they are hard.

A second question concerns the informational constraints on the model. We have assumed that the arrival intensity measure is known, and available as input to determining a socially optimal allocation or setting admission prices. In practice, this information is unlikely to be available, but needs to be inferred from observation. If a customer's delay sensitivity is revealed upon arrival, then the arrival distribution can easily be measured. But eliciting delay sensitivities truthfully can be a challenge in practice. It is an open question whether it is still possible to set admission prices in such a way as to ensure that the Wardrop equilibrium either coincides with the welfare optimizing allocation, or approximates it to within some factor.

Finally, we have assumed that a benevolent mechanism designer sets admission prices to maximize social welfare; it is interesting to ask what happens if the admission prices are set by a revenue maximizing service provider. Further, in such a revenue maximizing scenario it would be interesting to see if competing service providers can sustain differentiated services.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Acknowledgements

This work was done while the first author was affiliated with IIT Bombay, India. Bodas and Manjunath acknowledge support from the Bharti Centre for Communication at IIT Bombay, CEFIPRA and IFCAM.

Open Access

This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate

if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Received: 13 April 2021 Accepted: 14 July 2021

Published online: 13 September 2021

References

1. Acemoglu D, Makhdoumi A, Malekian A, Ozdaglar A (2018) Informational Braess' paradox: The effect of information on traffic congestion. *Oper Res* 66(4), 893–917
2. Afeche P, Mendelson H (2004) Pricing and priority auctions in queueing systems with a generalized delay cost structure. *Manage Sci* 50(7), 869–882
3. Balachandran KR (1972) Purchasing priorities in queues. *Manage Sci* 18(5), 319–326
4. Beckmann M, McGuire CB, Winsten CB (1956) *Studies in the economics of transportation*. Yale University Press
5. Bell CH, Stidham S (1983) Individual versus social optimization in the allocation of customers to alternative servers. *Manage Sci* 29(7):831–839
6. Borkar VS, Manjunath D (2004) Charge-based control of DiffServ-like queues. *Automatica* 40(12):2043–2057
7. Borst SC (1995) Optimal probabilistic allocation of customer types to servers. *ACM SIGMETRICS Perf Eval Rev* 23(1):116–125
8. Bradford RM (1996) Incentive compatible pricing and routing policies in multiserver queues. *Eur J Oper Res* 89:226–236
9. Dube P, Borkar VS, Manjunath D (2002) Differential join prices for parallel queues: social optimality, dynamic pricing algorithms and application to Internet pricing. In *Proc. IEEE INFOCOM*, pp 276–283
10. Edelson N, Hilderbrand D (1975) Congestion toll for Poisson queuing processes. *Econometrica* 43(1):81–92
11. Fleischer L, Jain K, Mahdian M (2004) Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In: *Proc. IEEE Symp. FOCS*, pp 277–285
12. Ganesh A, Laevens K, Steinberg R (2007) Congestion pricing and noncooperative games in communication networks. *Oper Res* 55(3):430–438
13. Hassin R, Haviv M (2003) *To queue or not to queue*. Kluwer Academic Publishers
14. Haviv M, Roughgarden T (2007) The price of anarchy in an exponential multi-server. *Oper Res Lett* 35(4):421–426

15. He Y, Magnac T (2018) A Pigouvian approach to congestion in matching markets. IZA Discussion Paper No. 11967, Institute of Labor Economics (IZA), Bonn
16. Jain R, Mullen T, Hausman R (2001) Analysis of Paris Metro pricing for QoS with a single service provider. In: Proc. IWQoS, Lecture Notes in Computer Science, Volume 2092/2001, pp 44–58
17. Kittsteiner T, Moldovanu B (2005) Priority auctions and queue disciplines that depend on processing time. *Manag Sci* 51(2), 236–248
18. Littlechild S (1974) Optimal arrival rate in a simple queueing system. *Int J Prod Res* 12(3):371–397
19. Lui F (1985) An equilibrium queueing model of bribery. *J Polit Econ* 93(4):760–781
20. Mankiw NG (2009) Smart taxes: an open invitation to join the Pigou club. *East Econ J* 35(1):14–23
21. Monderer D, Shapley LS (1996) Potential games. *Games Econ Behav* 14(1), 124–143
22. Naor P (1969) The regulation of queue size by levying tolls. *Econometrica* 37(1):15–24
23. Nisan N, Roughgarden T, Tardos E, Vazirani V (2007) *Algorithmic game theory*. Cambridge University Press
24. Odlyzko A (1999) Paris Metro pricing for the Internet. In: *Proceedings of the 1st ACM Conference on Electronic Commerce*, pp 140–147
25. Pigou AC (1920) *The economics of welfare*. Macmillan, London
26. Rao S, Petersen E (1998) Optimal pricing of priority services. *Oper Res* 46(1):46–56
27. Rosenthal R (1973) A class of games possessing pure-strategy Nash equilibria. *Int J Game Theory* 2(1):65–67
28. Roughgarden T (2005) *Selfish routing and the price of anarchy*. The MIT Press
29. Sethuraman J, Squillante M (1999) Optimal stochastic scheduling in multiclass parallel queues. *ACM SIGMETRICS Perform Evaluation Rev* 27(1):93–102
30. Smith MJ (1979) The marginal cost taxation of a transportation network. *Trans Res Ser B* 13B:237–242
31. Tandra R, Hemachandra N, Manjunath D (2004) Join minimum cost queue for multiclass customers: stability and performance bounds. *Probab Eng Inform Sci* 18(4):445–472
32. Wardrop JG (1952) Some theoretical aspects of road traffic research. *Proc Inst Civil Eng* 1(3):325–378
33. Yang H, Huang H-J (1998) Principle of marginal cost pricing: How does it work in a general road network? *Trans Res Ser A* 32(1), 45–54
34. Yang H, Huang H-J (2004) The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Trans Res Part B: Methodol* 38(1), 1–15



Tejas Bodas holds MTech and PhD degrees from IIT Bombay. He has been a visiting fellow at TIFR, Mumbai, a postdoc at LAAS-CNRS, Toulouse, a visiting postdoc at the University of Antwerp, and a C.V. Raman postdoc at IISc, Bangalore. He is currently an Assistant Professor at the

Electrical Engineering Department of IIT Dharwad. His research interests are in queuing theory, game theory, MDPs and reinforcement learning.



Ayalvadi Ganesh received his BTech from IIT Madras, and MS and PhD from Cornell University. He is currently an Associate Professor at the School of Mathematics, University of Bristol. He has previously worked at the University of Edinburgh, Birkbeck College,

Hewlett-Packard Labs and Microsoft Research. His research interests include large deviations, queuing theory, random graphs and distributed algorithms. He is a recipient of best publication awards from INFORMS (2005) and ACM SIGMETRICS (2010). He is co-author of the monograph *Big Queues* (2004).



D. Manjunath received his BE from Mysore University, MS from IIT Madras and PhD from Rensselaer Polytechnic Institute. He has been with the Electrical Engineering Department of IIT Bombay since July 1998 where he is now an Institute Chair Professor. He has previously

worked at the Corporate R&D Center of General Electric, the University of Delaware, the University of Toronto, and IIT Kanpur. His recent research has concentrated on random networks with applications in wireless and sensor networks, network pricing and queue control. He is a recipient of the best paper award at ACM SIGMETRICS 2010. He was an associate editor of *IEEE Transactions on Networking, Queueing Systems: Theory and Applications*, and *Sadhana: The Proceedings of the Indian Academy of Sciences*. He is a coauthor of two textbooks, *Communication Networking: An Analytical Approach* (May 2004) and *Wireless Networking* (Apr 2008).