

A Study on Price-Discrimination for Robust Social Coordination

Philip N. Brown and Jason R. Marden

Abstract—In many modern settings, system-level efficiency depends greatly on the choices of the system’s users. For example, congestion in traffic networks is often the result of self-interested choices made by individual drivers or users. An important research focus involves characterizing road tolls that bring users’ incentives in line with the system planner’s goals. Recent results have proposed methods of computing tolls that rely on very little instance-specific information, e.g. network topology or users’ value-of-time. Unfortunately, these existing methodologies require very large tolls, which may be impractical for a number of reasons. In this paper, we study price-discrimination as a possible means to social coordination; that is, we propose a model in which a system planner has a limited ability to disaggregate a user population into a number of groups on the basis of users’ value-of-time, and charge different tolls to different groups. We show that even without access to problem-specific information, it is possible to incentivize arbitrarily-low-congestion flows by discriminating finely enough. Second, we provide a general method of adapting standard taxation techniques to discriminatory settings and we analytically characterize the resulting efficiency gains. Finally, we look specifically at parallel-network, affine-cost routing games, and show how to derive the congestion-minimizing discriminatory taxation functions and prove upper-bounds on the inefficiencies resulting from this method.

I. INTRODUCTION

As engineered systems become more tightly connected with their end-users, system performance becomes increasingly dependent on individual users’ behavior. It is a widely-studied fact that self-interested behavior by users can severely degrade system-level performance [1]–[3]. As a result, an emerging engineering challenge is the problem of influencing social behavior (via incentives, information, or otherwise) to improve system performance [4]–[9].

One setting of interest is that of distributed network routing problems, in which a unit of traffic must be routed from some source to some destination across a network with the goal of minimizing the average travel time of the traffic. In general, if the traffic is composed of a large collection of independent agents each making individual routing decisions, it is a well-studied fact that the emergent social behavior can have extremely poor efficiency [1]. A common approach for mitigating this inefficiency is to charge specially-designed taxes on over-congested roads, hoping to influence the agents to route more efficiently [11].

This research was supported by AFOSR grant #FA9550-12-1-0359, ONR grant #N00014-12-1-0643, and NSF Grant #ECCS-1351866.

P. N. Brown is a graduate research assistant with the Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO 80309, philip.brown@colorado.edu. Corresponding author.

J. R. Marden is with the Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO 80309, jason.marden@colorado.edu.

A challenge in this setting is that an incentive-designer may not have access to a perfect system characterization at design-time; for example, the exact distribution of network users’ toll-sensitivities may not be known, or the designer may desire a degree of robustness to sudden changes in topology caused by natural disasters. It is thus crucial to understand the informational dependencies of any incentive-design methodology. A few examples are: fixed tolls (constant functions of flow on each link) can incentivize any desired network flow, but require copious amounts of information [11], [12]; marginal-cost tolls (which are strictly flow-varying) can incentivize optimal flows without requiring information on network topology [10], [13]; unbounded so-called “universal” tolls can incentivize optimal flows without requiring any instance-specific information [14].

The above results all charge every agent on each edge the same price. It is almost trivial to see that if the tax-designer has the ability to charge each agent an individualized price, he would be able to apply an individualized marginal-cost pricing scheme and inherit the optimality given by [13] while avoiding the unbounded tolls of [14]. This individualized pricing approach is suggested in [15].

This concept of perfect price-discrimination (also called first-degree price discrimination [?]) is theoretically appealing, but likely to be impractical in a real-world setting. However, it may be possible to perform price-discrimination at a coarser level; that is, to partition the users into a small number of groups or “bins” on the basis of their price-sensitivity and charge a single price to members of each bin. It seems natural to assume that even coarse price-discrimination could dramatically improve the efficiency guarantees, since each agent would be charged a price that is close to the “correct” price for his particular price-sensitivity.

The concept of classifying agents according to their preferences is not new, and has been studied extensively by economists for purposes of revenue-maximization. If a seller lacks direct access to customers’ sensitivities, it may still be possible to indirectly disaggregate the customer population: senior, student, and corporate discounts are common means to this end [?]. Even in cases such as public utility pricing and internet service providers in which certain kinds of discriminatory pricing are prohibited by law, price-discrimination via volume discounts (known as “nonlinear” pricing) is a common practice [22]. Price discrimination has also been studied in the context of cloud computing [17] and the provision of network services [18].

In this paper, we adopt a simple model of price-discrimination in order to study its possible benefits in social coordination. In our model, the tax-designer has the ability

to group network users into “bins” according to their price-sensitivity and charge all members of each bin a single price. Hopefully, this results in each user paying a price “close” to that user’s optimal price. We assume that the designer has no distributional information about users *within* each bin (similar to an approach outlined in [19]), but that each user is correctly categorized, however coarsely.

First, we demonstrate in Theorem 1 that network flows can be made arbitrarily close to optimal flows by binning users sufficiently finely. This validates the intuitive concept that if we charge each user an individualized price, we can enforce any network flow we desire.

However, Theorem 1 gives no hint as to how many bins are required to achieve a particular efficiency target; accordingly, in Theorem 2, we prove a fundamental equivalence between discriminatory pricing and the tax-designer’s uncertainty regarding the user population’s price-sensitivity. Here, we show that discriminatory pricing on a poorly-characterized population is essentially identical to simple pricing on a well-characterized population.

Finally, in Theorem 3, we consider the class of affine-cost, parallel-network routing games and provide a methodology for deriving the optimal binning and bin taxes for any number of bins. We also prove bounds on the inefficiencies resulting from this congestion-minimizing binning.

II. MODEL AND SUMMARY OF CONTRIBUTIONS

A. Routing Game

Consider a network routing problem in which a unit mass of traffic needs to be routed across a network (V, E) , which consists of a vertex set V and edge set $E \subseteq (V \times V)$. We call a source/terminus vertex pair $(s_c, t_c) \in (V \times V)$ a *commodity*, and the set of all commodities \mathcal{C} . We assume that for each $c \in \mathcal{C}$, there is a mass of traffic $r_c > 0$ that needs to be routed from s_c to t_c . We write $\mathcal{P}_c \subset 2^E$ to denote the set of *paths* available to traffic in commodity c , where each path $p \in \mathcal{P}_c$ consists of a set of edges connecting s_c to t_c . Let $\mathcal{P} = \cup \{\mathcal{P}_c\}$.

A *feasible flow* $f \in \mathbb{R}^{|\mathcal{P}|}$ is an assignment of traffic to various paths such that for each commodity, $\sum_{p \in \mathcal{P}_c} f_p = r_c$, where $f_p \geq 0$ denotes the mass of traffic on path p . Without loss of generality, we assume that $\sum_{c \in \mathcal{C}} r_c = 1$.

Given a flow f , the flow on edge e is given by $f_e = \sum_{p: e \in p} f_p$. To characterize transit delay as a function of traffic flow, each edge $e \in E$ is associated with a specific latency function $\ell_e : [0, 1] \rightarrow [0, \infty)$. We adopt the standard assumptions that latency functions are nondecreasing, continuously differentiable, and convex. We measure the efficiency of a flow f by its *total latency*, given by

$$\mathcal{L}(f) = \sum_{e \in E} f_e \cdot \ell_e(f_e) = \sum_{p \in \mathcal{P}} f_p \cdot \ell_p(f_p), \quad (1)$$

where $\ell_p(f) = \sum_{e \in p} \ell_e(f_e)$ denotes the latency on path p . We denote the flow that minimizes the total latency by

$$f^* \in \underset{f \text{ is feasible}}{\operatorname{argmin}} \mathcal{L}(f). \quad (2)$$

Due to the convexity of ℓ_e , $\mathcal{L}(f^*)$ is unique.

A *routing problem* is specified by the tuple $G = (V, E, \mathcal{C}, \{\ell_e\})$. We write the set of all such routing problems as \mathcal{G} . We will often use shorthand notation such as $e \in \mathcal{G}$ to denote $(e \in G : G \in \mathcal{G})$.

To study taxation mechanisms for influencing the emergent collective behavior resulting from self-interested price-sensitive users, we model the above routing problem as a non-atomic game. We assign each edge $e \in E$ a flow-dependent, nondecreasing taxation function $\tau_e : [0, 1] \rightarrow \mathbb{R}^+$. We characterize the taxation sensitivities of the users in commodity c with a monotone, nondecreasing function $s^c : [0, r_c] \rightarrow [S_L, S_U]$, where each user $x \in [0, r_c]$ has a taxation sensitivity $s_x^c \in [S_L, S_U] \subseteq \mathbb{R}^+$ and $S_U \geq S_L \geq 0$ denote upper and lower sensitivity bounds, respectively. Given a flow f , the cost that user x experiences for using path $\tilde{p} \in \mathcal{P}_c$ is of the form

$$J_x(f) = \sum_{e \in \tilde{p}} [\ell_e(f_e) + s_x^c \tau_e(f_e)], \quad (3)$$

and we assume that each user prefers the lowest-cost path from the available source-destination paths.

In our analysis, we assume that each sensitivity distribution function s^c is unknown; for a given routing problem G and $S_U \geq S_L \geq 0$ we define the set of possible sensitivity distributions as the set of monotone, nondecreasing functions $\mathcal{S}_G = \{s^c : [0, r_c] \rightarrow [S_L, S_U]\}_{c \in \mathcal{C}}$. We write $s \in \mathcal{S}_G$ to denote a specific collection of sensitivity distributions, which we term a *user population*.

In this paper, we study the effects of price-discrimination; i.e., the practice of charging different prices to different users. To model this, we assume that we have some ability to group (or “bin”) users according to their price-sensitivity, and that we can charge different tolls to different bins. Such a price-discrimination scheme is comprised of two components: a collection of bins represented by a partition of the interval $[S_L, S_U]$ into m sub-intervals, and a collection of taxation functions for each group.

We write the collection of bin boundaries as $\{\beta_i\}_{i=0}^m$, $\beta_0 = S_L$, $\beta_m = S_U$, with $\beta_i < \beta_{i+1} \forall i < m$. For each edge e , we charge all users in bin i (that is, all users whose sensitivities lie in the interval $[\beta_{i-1}, \beta_i]$) a taxation function $\tau_e^i(f_e)$, yielding for each edge e a collection of m distinct taxation functions $\{\tau_e^i\}_{i=1}^m$. So that we can make general statements about binnings as functions of m , we write B to represent a function that maps each $m \in \mathbb{N}$ to a particular partition $\{\beta_i\}_{i=0}^m$ and taxation functions $\{\{\tau_e^i\}_{e \in E}\}_{i=1}^m$, and write B^m to denote a binning for a specific value of m . Occasionally, we use the notation B^1 to denote a trivial “binning” in which all users are charged the same price.

We model the social behavior resulting from an m -binning as a Nash flow, or a flow f in which for all users x , where x belongs to commodity $c \in \mathcal{C}$ and bin i , we have

$$J_x(f) = \min_{p \in \mathcal{P}_c} \left\{ \sum_{e \in p} [\ell_e(f_e) + s_x^c \tau_e^i(f_e)] \right\}. \quad (4)$$

It is well-known that a Nash flow exists for any non-atomic game of the above form [20]; further, such Nash flows are essentially unique [13].

For a given routing problem $G \in \mathcal{G}$, we gauge the efficacy of a binning B^m by comparing the total latency of the resulting Nash flow and the total latency associated with the optimal flow, and then performing a worst-case analysis over all possible sensitivity distributions. Let $\mathcal{L}^*(G)$ denote the total latency associated with the optimal flow, and $\mathcal{L}^{\text{nf}}(G, s, B^m)$ denote the total latency of the Nash flow resulting from binning B^m and user population s . The worst-case efficiency loss associated with this specific instance is captured by the *price of anarchy* which takes the general form

$$\text{PoA}(G, S_L, S_U, B^m) = \sup_{s \in \mathcal{S}_G} \left\{ \frac{\mathcal{L}^{\text{nf}}(G, s, B^m)}{\mathcal{L}^*(G)} \right\} \geq 1. \quad (5)$$

B. Summary of Our Contributions

An overview of our contributions is as follows:

Theorem 1:

We define a family of binnings which guarantee optimal Nash flows for any routing problem, provided we discriminate finely enough.

Theorem 2:

We show that for any routing problem, the price of anarchy resulting from discriminatory pricing is no worse than the price of anarchy resulting from non-discriminatory pricing for a better-characterized population.

Theorem 3:

We derive the *optimal* discriminatory pricing for parallel-network, affine-cost routing problems and affine taxation mechanisms.

Our first result in Theorem 1 is to show that in general, we can enforce arbitrarily-efficient Nash flows by using a large number of bins with no instance-specific information; i.e., network topology, user sensitivities, or user demands. This is similar to the approach outlined in [14], in which large tolls (again requiring no instance-specific information) are used to enforce highly-efficient Nash flows. However, it should be noted that the methods of [14] require arbitrarily-large tolls to enforce optimal flows; in contrast, our discriminatory pricing can always enforce optimal flows with bounded tolls.

Our second contribution in Theorem 2 is a general characterization result showing that discriminatory pricing on a poorly-characterized population is in some sense equivalent to non-discriminatory pricing on a well-characterized population. That is, discriminatory pricing effectively reduces the designer’s uncertainty regarding the users’ price-sensitivity. This presents the incentive-designer with a powerful tool for implementing discriminated prices, effectively allowing a designer to drag-and-drop an undiscriminated pricing design directly into a discriminated one.

In Theorem 3, we consider these concepts in the context of a question posed in [14]: “in affine-cost, parallel-network congestion games, if edge tolls are never allowed to exceed some upper bound T , what prices are optimal?” Here, given

an upper-bound on edge tolls, we ask what bin boundaries and tolls minimize the price of anarchy. We investigate the fundamental equivalence between fine discrimination and large tolls and analytically characterize the complementary benefits of the two approaches. We show how to compute the optimal bin boundaries, and how to charge tolls in each bin to minimize congestion. Finally, we derive an upper bound on the price of anarchy resulting from these prices. See Figure 1 for a depiction of the difference between the generic approach given in Theorem 2 and the specific, congestion-minimizing approach from Theorem 3.

III. RELATED WORK

There has been significant research geared towards developing taxation mechanisms to eradicate the inefficiency caused by users’ self-interested routing choices; however, to our knowledge, price-discrimination has not been studied for this purpose in congestion games. Here, we survey some relevant results pertaining to pricing for congestion-minimization in network routing, highlighting in particular the tradeoffs between the sophistication required of a taxation mechanism and its corresponding informational requirements.

– *Fixed tolls: low sophistication, high information requirement:* These taxation mechanisms assign tolls that are constant functions of flow (i.e., for any $e \in G$, $\tau_e(f_e) = q_e$ for some $q_e \geq 0$). It is known that if a system-planner has access to complete information about a routing problem (e.g., network topology, user demands, exact distribution of user sensitivities), fixed tolls can be computed which enforce any feasible network flow [11], [12]. Computing fixed tolls typically requires the system-planner to possess the above information, and the robustness of these tolls to mischaracterizations of network topology is an open question.

– *Marginal-Cost Tolls: medium sophistication, medium informational requirement:* Also called “Pigovian” or “congestion” tolls, these guarantee optimal flows without the need for information regarding network topology, but require edge taxation functions to be strictly flow-varying. These tolls are of the following form: for any $e \in \mathcal{G}$ with latency function ℓ_e , the accompanying taxation function is

$$\tau^{\text{mc}}(f_e) = f_e \cdot \frac{d}{df_e} \ell_e(f_e), \quad \forall f_e \geq 0. \quad (6)$$

In [13] it is shown that for any $G \in \mathcal{G}$ we have $\mathcal{L}^*(G) = \mathcal{L}^{\text{nf}}(G, s, \tau^{\text{mc}})$ provided that all users have a sensitivity exactly equal to 1. The authors of [4] show that if the user population is heterogeneous in price-sensitivity, marginal-cost taxes scaled by $\sqrt{S_L S_U}$ do possess a degree of robustness to mischaracterizations of user sensitivities, but can no longer guarantee a price of anarchy of 1.

– *“Universal” Tolls: High sophistication, low informational requirement:* In [14], it is shown that if a designer possesses no information regarding network topology, user sensitivities, or user demands, optimal network flows can be enforced – *if and only if* the tax-designer levies arbitrarily-large tolls.

In contrast to the network-routing literature, the economic literature on price-discrimination is concerned largely with profit-maximization. It is well-known that a monopolist can

maximize profits with “first-degree” (also called “perfect”) price-discrimination, in which prices are individualized for every customer [?], provided that goods cannot be re-sold. Another commonly-studied form of price-discrimination is often termed “second-degree” price discrimination or nonlinear pricing, in which different prices are charged for different quantities of a good or service [21], essentially inducing customers to partition themselves into different sensitivity classes. In this paper, we study a coarse version of first-degree discrimination, under the common assumption that perfect price-discrimination is difficult in practice.

IV. OUR CONTRIBUTIONS

A. Universal Discriminatory Pricing

As discussed in the previous section, the state of the art gives us no obvious way to enforce optimal flows without either significant quantities of information or excessively-high tolls. In this paper, we show that price-discrimination may provide a third way. We begin with our most general result, in which we present a family of binnings which enforce arbitrarily-optimal flows for any routing problem. This optimality is asymptotic in the number of bins, which implies that it is always possible to enforce Nash flows within ϵ of optimal with a finite number of bins.

Theorem 1: Define \mathcal{B} as the family of binnings whose bin sizes shrink to 0 as m approaches infinity; for any $B^m \in \mathcal{B}$,¹

$$\lim_{m \rightarrow \infty} (\beta_i^m - \beta_{i-1}^m) = 0, \quad (7)$$

and the bin taxation functions of $B^m \in \mathcal{B}$ satisfy the following: for each bin i choose any $\kappa_i^m \in \left[\frac{1}{\beta_i^m}, \frac{1}{\beta_{i-1}^m}\right]$ and let edge tolls be given by

$$\tau_e^i(f_e) = \kappa_i^m f_e \cdot \frac{d}{df_e} \ell_e(f_e). \quad (8)$$

Then for any $G \in \mathcal{G}$ and any binning $B^m \in \mathcal{B}$,

$$\lim_{m \rightarrow \infty} \text{PoA}(G, S_L, S_U, B^m) = 1 \quad (9)$$

and for all m , all tolling functions are bounded by $\max_e \ell_e'(1)/S_L$.

Note that B^m need not depend on information about user sensitivity distributions or demands, and may be completely topology-independent.

Proof: For each $m \in \mathbb{N}$, let binning B^m be defined with bin boundaries according to (7) and bin tolls according to (8), so that $B^m \in \mathcal{B}$. For any routing problem $G \in \mathcal{G}$ and price-sensitivities $s \in \mathcal{S}$, let $f^m = (f_p^m)_{p \in \mathcal{P}}$ denote the Nash flow resulting from the binning B^m . For each commodity c , let $\mathcal{P}_c^m \subseteq \mathcal{P}_c$ denote the set of paths that have positive flow in f^m . For any $p \in \mathcal{P}_c^m$, there must be some user $x \in [0, r_c]$ using p ; suppose this user has sensitivity $s_x^c \in [\beta_{i-1}^m, \beta_i^m]$, then the cost experienced by this user is given by

$$J_x(f^m) = \sum_{e \in p} \left[\ell_e(f_e) + s_x^c \kappa_i^m f_e \cdot \frac{d}{df_e} \ell_e(f_e) \right].$$

¹For the sake of precision, in this theorem we make explicit the dependence of each bin boundary on m by writing β_i^m .

Write $\ell_e^*(f_e) = f_e \cdot \frac{d}{df_e} \ell_e(f_e)$; then for any other path $p' \in \mathcal{P}_c \setminus p$, user x must experience a lower cost on p than on p' , or

$$\sum_{e \in p} \ell_e(f_e) - \sum_{e \in p'} \ell_e(f_e) \leq s_x^c \kappa_i^m \left[\sum_{e \in p'} \ell_e^*(f_e) - \sum_{e \in p} \ell_e^*(f_e) \right]. \quad (10)$$

Therefore, for any $m \geq 1$, f^m must satisfy some set of inequalities defined by (10). By definition,

$$\beta_{i-1}^m / \beta_i^m \leq s_x^c \kappa_i^m \leq \beta_i^m / \beta_{i-1}^m,$$

so (7) implies that $\lim_{m \rightarrow \infty} s_x^c \kappa_i^m = 1$. Thus, because all functions in (10) are continuous, f^m converges to a set F^* of feasible flows that satisfy

$$\sum_{e \in p} \ell_e(f_e) - \sum_{e \in p'} \ell_e(f_e) \leq \left[\sum_{e \in p'} \ell_e^*(f_e) - \sum_{e \in p} \ell_e^*(f_e) \right] \quad (11)$$

for all c , all $p \in \mathcal{P}_c^*$, and $p' \in \mathcal{P}_c$, where $\mathcal{P}_c^* \subseteq \mathcal{P}_c$ is some subset of paths. Inequalities (11) (combined with the feasibility constraints on f) specify a Nash flow for G for a unit-sensitivity population with marginal-cost taxes as defined in (6); any such Nash flow must be optimal [13]; that is, any $f \in F^*$ is a minimum-latency flow for G . Thus, since $\mathcal{L}(f)$ is a continuous function of f ,

$$\lim_{m \rightarrow \infty} \mathcal{L}(f^m) = \mathcal{L}^*(G), \quad (12)$$

obtaining the proof of the theorem. \blacksquare

B. General Effect of Discriminatory Pricing

Theorem 1 showed that we can enforce low-congestion routing using price-discrimination, but it gave no hint as to how many bins are needed or how the price of anarchy evolves as a function of m . This is the purview of Theorem 2, in which we show that there is a general equivalence between fine discrimination and well-characterized sensitivity distributions. Here, we show that the price of anarchy resulting from discriminatory pricing for a poorly-characterized population is no worse than the price of anarchy resulting from non-discriminatory pricing for a well-characterized population.

Theorem 2: Suppose for routing problem G that some taxation mechanism $\tau(S_L, S_U)$ is known to have price of anarchy $\text{PoA}(G, S_L, S_U)$. For any $S'_L > 0$, let $S'_U = S'_L (S_U/S_L)^{1/m}$, and define the bin boundaries of B^m by

$$\beta_i = S'_L \frac{m-i}{m} S'_U \frac{i}{m}. \quad (13)$$

Then if the bin taxes of B^m are given by $\tau_i = S'_L / \beta_{i-1} \cdot \tau(S'_L, S'_U)$, the following holds:

$$\text{PoA}(G, S_L, S_U, B^m) \leq \text{PoA} \left(G, S'_L, S'_L \left(\frac{S_U}{S_L} \right)^{1/m}, B^1 \right). \quad (14)$$

In particular, we wish to point out two important facts regarding Theorem 2. First, it is natural to evaluate the uncertainty of our user-sensitivity estimate by the sensitivity ratio

S_U/S_L : the higher the ratio, the less certainty we possess. Theorem 2 shows us that m -binning reduces the effective sensitivity ratio to $(S_U/S_L)^{1/m}$. Thus, by applying price discrimination, we can dramatically reduce our uncertainty regarding the price-sensitivity of network users.

Second, note that the guarantees of Theorem 2 are independent of the specific taxation mechanism used; thus, this result can be used as a design tool to take any off-the-shelf taxation mechanism and apply it to a discriminatory setting.

However, it is important to understand that the price of anarchy provided by Theorem 2 need not be *optimal* in any sense. If a binning is designed more precisely with a particular taxation methodology in mind, it may be possible to guarantee even better network efficiencies. This is the focus of Theorem 3, in which we look at a restricted class of routing problems and taxation mechanisms and derive the optimal binning in that specific setting.

Proof: Consider routing problem G and population s ; design bin boundaries and charge taxes according to the theorem statement. Note that for every i , (13) implies that

$$\frac{\beta_i}{\beta_{i-1}} = \frac{S'_U}{S'_L} = \left(\frac{S_U}{S_L} \right)^{1/m}; \quad (15)$$

that is, each bin has the same “width” as the entire emulated population $s' \in [S'_L, S'_U]$. Consider the cost experienced on any edge e by some agent x who happens to belong to bin i :

$$J_x(f) = \ell_e(f_e) + \kappa_i s_x \tau_e(f_e). \quad (16)$$

Inserting the definition of κ_i in the above, we obtain

$$J_x(f) = \ell_e(f_e) + \frac{s_x S'_L}{\beta_{i-1}} \tau_e(f_e). \quad (17)$$

But $s_x \in [\beta_{i-1}, \beta_i]$ implies that

$$\ell_e(f_e) + S'_L \tau_e(f_e) \leq J_x(f) \leq \ell_e(f_e) + S'_U \tau_e(f_e), \quad (18)$$

or that agent x sees the exact same cost as he would if he were a member of population s' . Since this is true of every agent for every edge in every bin, it must be true that every Nash flow resulting from these binned tolls corresponds exactly to a Nash flow resulting from the nominal tolls $\{\tau\}$ and some sensitivity distribution in $[S'_L, S'_U]$. In particular, no binned Nash flow can have higher total latency than the worst-case flows for populations in $[S'_L, S'_U]$. Since $S'_U/S'_L = (S_U/S_L)^{1/m}$, the theorem conclusion follows. ■

C. Optimal Binning for Simple Routing Problems

The principles proved in Theorem 2 are compelling, but in general may not result in optimal binning. That is, when investigating a restricted class of games, it may often be the case that we can design pricing that significantly outperforms the pricing described by Theorem 2.

In this section, we restrict attention to the class of parallel-network, affine-cost routing games. For the following results, let $\mathcal{G}^p \subseteq \mathcal{G}$ represent the class of all single-commodity,

parallel-link routing problems with affine latency functions. That is, for all $e \in \mathcal{G}^p$, the latency function satisfies

$$\ell_e(f_e) = a_e f_e + b_e \quad (19)$$

where a_e and b_e are non-negative edge-specific constants. “Single-commodity” implies that all traffic has access to all network edges. Furthermore, we assume that every edge has positive flow in an un-tolled Nash flow.²

Following from [14], we endeavor to apply optimal bounded tolls to this game which have no dependence on the specific network to which they are applied. We want to define a general rule for assigning edge taxation functions in which each edge’s taxes depend only on that edge’s congestion properties. Efficiency guarantees resulting from such a network-agnostic taxation mechanism are thus independent of specific network topologies, meaning that they are robust to edge deletion or sudden changes in topology.

Definition 1: The *Bounded Affine Taxation Mechanism* assigns tolls of

$$\tau_e(f_e) = \kappa_1 a_e f_e + \kappa_2 b_e, \quad (20)$$

where a_e and b_e are the latency function coefficients in (19) and $\kappa_1 \leq \kappa_{\max}$ and $\kappa_2 \leq \kappa_{\max}$ are non-negative edge-independent constants upper-bounded by some $\kappa_{\max} \geq 0$.

In [14], it is shown that for non-discriminatory pricing, this taxation mechanism is optimal over the space of all bounded network-agnostic taxation mechanisms. Thus, it seems a natural mechanism to study in our discriminatory setting.³

Now, in Theorem 3, we show how to compute optimal bin boundaries and affine taxation function coefficients that minimize congestion for any κ_{\max} . Furthermore, we derive an upper bound for the price of anarchy that is independent of the number of network links and holds for any S_U .

Theorem 3: For any $G \in \mathcal{G}^p$, for any set of bin boundaries $\{\beta_i\}_{i=1}^m$, the optimal bounded affine tolling coefficients are given by

$$\kappa_1^i = \kappa_{\max} \quad (21)$$

$$\kappa_2^i = \max \left\{ 0, \frac{(\kappa_1^i)^2 \beta_{i-1} \beta_i - 1}{\beta_{i-1} + \beta_i + 2(\kappa_1^i) \beta_{i-1} \beta_i} \right\}. \quad (22)$$

Furthermore, the congestion-minimizing bin boundaries $\{\beta_i\}_{i=1}^m$ can be found by solving Optimization Problem (P) (see Figure 2). Let $\lambda = \frac{\kappa_{\max} S_U}{1 + \kappa_{\max} S_L}$. If $\kappa_{\max} \leq 1/S_L$, let $\mu = S_L \kappa_{\max}$; otherwise, let $\mu = 1$. Then for all $S_U \in [S_L, \infty]$ and this binning B^m ,

$$\text{PoA}(\mathcal{G}^p, S_L, S_U, B^m) \leq \frac{4}{3} \left(1 - \frac{\min\{\lambda^{1/m}, \mu\}}{(1 + \min\{\lambda^{1/m}, \mu\})^2} \right). \quad (23)$$

²This is essentially a regularity condition that prevents the creation of unrealistic, highly-pathological networks; e.g., if a network contains an edge with a very high constant latency function, tolling functions could cause highly-sensitive users to divert to this edge, causing gross network “inefficiencies.” Note that we can always assign infinite tolls to such unused edges to ensure that the regularity condition is met.

³However, note that in this paper we do not prove the optimality of affine tolls over the space of all taxation mechanisms, we merely prove the optimal choices of κ_1 and κ_2 .

Optimization Problem (P)

$$\begin{aligned}
 & \text{Max}_{\{\beta_i\}_{i=1}^m} \gamma_L \\
 & \text{s.t. } \gamma_L \leq \frac{\kappa_{\max} + 1/\beta_i}{\kappa_{\max} + 1/\beta_{i-1}} \quad \forall i \in \{1, \dots, m\} \quad (24) \\
 & \quad \beta_{i-1} \leq \beta_i \quad \forall i \in \{1, \dots, m\} \\
 & \quad \beta_0 = S_L \\
 & \quad \beta_m = S_U \\
 & \quad \beta_1 \geq \frac{1}{\kappa_{\max}^2 S_L}. \quad (25)
 \end{aligned}$$

Fig. 1. As proved in Theorem 3, the solutions $\{\beta_i\}_{i=1}^m$ to this optimization problem are congestion-minimizing bin boundaries for any routing problem.

A few words are in order regarding the price of anarchy bound in (21): First, this bound is tight for cases when $S_U = \infty$; i.e., there is no upper bound on the price-sensitivities of the agents. When S_U is finite, the tools of Lemmas 3.1 and 3.2 can be used to determine an exact price of anarchy once the optimal bin boundaries have been derived. We are not aware of a convenient closed-form expression for the exact price of anarchy for finite S_U , but in Figure 1 we show that even for relatively low values of S_U , the gap between the finite- S_U and infinite- S_U prices of anarchy is not large.

Second, since we are dealing with bounded tolls, whenever $\kappa_{\max} < 1/S_L$, it is not possible to guarantee perfectly-optimal network flows, even for arbitrarily-high m . This is because when κ_{\max} is too low, a homogeneous sensitivity distribution with $s_x \equiv S_L$ cannot be effectively influenced, and after a point, finer binning cannot remedy this. This is captured in the theorem statement by the parameter μ , which prevents extremely-fine binning from improving the price of anarchy when κ_{\max} is too low.

See Figure 3 for a depiction of the congestion-minimizing bin boundaries for several values of m . In Figure 1, we depict the price of anarchy as a function of m , contrasting the guarantees provided by Theorems 2 and 3.

As a first step towards proving Theorem 3, we introduce Lemma 3.1, in which we present a powerful tool with which we can analyze the price of anarchy of parallel affine congestion games under various types of tolls.

Lemma 3.1: Suppose that there exists a function $\gamma : [0, 1] \rightarrow [\gamma_L, \gamma_U]$ with $0 \leq \gamma_L \leq 1$ and $\gamma_U \leq 1/\gamma_L$ such that in every $G \in \mathcal{G}^p$, the cost function of user x on edge e is given by

$$J_x^e(f_e) = (1 + \gamma(x)) a_e f_e + b_e. \quad (26)$$

Then the price of anarchy of \mathcal{G}^p is tightly upper-bounded by

$$\text{PoA}(\mathcal{G}^p) \leq \frac{4}{3} \left(1 - \frac{\gamma_L}{(1 + \gamma_L)^2} \right). \quad (27)$$

Proof: The simplest proof of this relates (24) to the analysis of scaled marginal-cost tolls presented in [4]. There, in Lemma 1.2, the authors show that for $\kappa \leq 1/\sqrt{S_L S_U}$, of all the Nash flows induced by edge tolls $\tau_e(f_e) = \kappa a_e f_e$,




Fig. 2. Optimal bin boundary locations computed by Optimization Problem (P) for $m \in \{1, \dots, 10\}$, with $\kappa_{\max} = S_L = 1$ and $S_U = 10$. These correspond exactly to the price of anarchy plot shown in Figure 1.

the worst congestion always occurs for a homogeneous population in which all users' sensitivities are equal to S_L . To prove Lemma 3.1, we shall compute a virtual sensitivity distribution s^v and tolling coefficient κ^v which will induce Nash flows that precisely mimic the behavior of Nash flows induced by γ (that is, Nash flows for cost functions (24)).

Given the γ function of the statement of Lemma 3.1, let $\kappa^v = 1$, and let $s_x^v = \gamma(x)$ for all $x \in [0, 1]$. By this definition, any Nash flow induced by γ has a corresponding marginal-cost-tolled Nash flow induced by a sensitivity distribution given by s^v ; we can accordingly use marginal-cost toll arguments to argue about γ -induced Nash flows. The upper and lower bounds of our virtual sensitivity distribution are thus given by $S_L^v = \gamma_L$ and $S_U^v = \gamma_U$, respectively. By the properties of γ_L and γ_U , it is always true that $\kappa^v \leq 1/\sqrt{S_L^v S_U^v}$, so the congestion-maximal Nash flows occur for a virtual homogeneous population with all users' sensitivities equal to S_L^v .

This implies that a Nash flow with all cost functions equal to $J_x^e(f_e) = (1 + \gamma_L) a_e f_e + b_e$ will have higher total congestion than a Nash flow induced by γ , and we can use techniques from [4] (namely, Lemma 1.3) to derive (25). The tightness of the bound in (25) is due to the constructive nature of the proofs from [4]. ■

Next, in Lemma 3.2, for any arbitrary bin boundaries, we derive specific tolling coefficients which should be charged in each bin.

Lemma 3.2: For any $G \in \mathcal{G}^p$, for any bin boundaries $\{\beta_i\}$, the optimal tolling coefficients can be obtained for each bin i by choosing

$$\kappa_1^i = \kappa_{\max} \quad (28)$$

$$\kappa_2^i = \max \left\{ 0, \frac{(\kappa_1^i)^2 \beta_{i-1} \beta_i - 1}{\beta_{i-1} + \beta_i + 2(\kappa_1^i) \beta_{i-1} \beta_i} \right\}. \quad (29)$$

With these coefficients, in the language of Lemma 3.1, it holds that

$$\gamma_L = \min_i \left\{ \min \left\{ \beta_{i-1} \kappa_{\max}, \frac{\kappa_{\max} + 1/\beta_i}{\kappa_{\max} + 1/\beta_{i-1}} \right\} \right\}. \quad (30)$$

Proof: Since uniform scaling by a constant factor does not change underlying Nash flows, we can say without loss of generality that the effective cost to agent $x \in [\beta_{i-1}, \beta_i]$

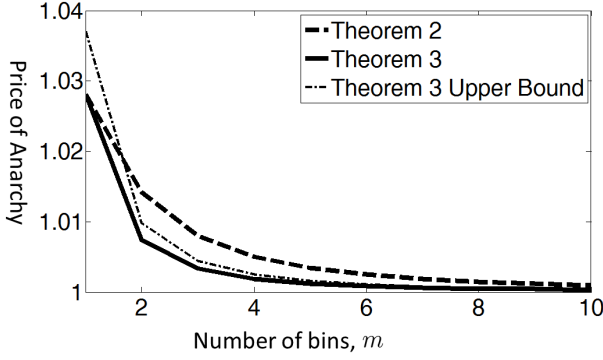


Fig. 3. Comparison between the price of anarchy resulting from the generic price-discrimination approach of Theorem 2 (dashed line), the specific congestion-minimizing price-discrimination approach of Theorem 3 (solid line), and the general upper-bound given in Theorem 3 (dash-dot line). For the first two curves, we apply affine tolls to the class of parallel-network, affine-cost congestion games, with $\kappa_{\max} = S_L = 1$ and $S_U = 10$. The dash-dot line represents the general upper-bound proved in Theorem 3 that holds for this value of κ_{\max} and S_L , and any S_U . Note how close this upper bound is to the instance-specific solid line for $m > 1$.

in bin i for edge e is given by

$$J_x(f_e) = \frac{1 + \kappa_1^i s_x}{1 + \kappa_2^i s_x} a_e f_e + b_e, \quad (31)$$

and when $\kappa_1^i \geq \kappa_2^i$, it is evident that⁴

$$\frac{1 + \kappa_1^i \beta_{i-1}}{1 + \kappa_2^i \beta_{i-1}} a_e f_e + b_e \leq J_x(f_e) \leq \frac{1 + \kappa_1^i \beta_i}{1 + \kappa_2^i \beta_i} a_e f_e + b_e. \quad (32)$$

For each i , for each $x \in [\beta_{i-1}, \beta_i]$, define $\gamma(x) \triangleq \frac{s_x(\kappa_1^i - \kappa_2^i)}{1 + \kappa_2^i s_x}$. Then any agent $x \in [0, 1]$ has the following cost for edge e :

$$J_x^e(f_e) = (1 + \gamma(x)) a_e f_e + b_e, \quad (33)$$

just as in (24). It is evident that in accordance with the assumptions of Lemma 3.1 we have

$$\begin{aligned} \gamma(x) &\geq \gamma_L \triangleq \min_i \{ \beta_{i-1} (\kappa_1^i - \kappa_2^i) / (1 + \kappa_2^i \beta_{i-1}) \} \text{ and} \\ \gamma(x) &\leq \gamma_U \triangleq \max_i \{ \beta_i (\kappa_1^i - \kappa_2^i) / (1 + \kappa_2^i \beta_i) \} \end{aligned} \quad (34)$$

Equation (25) in Lemma 3.1 shows that we minimize the price of anarchy by maximizing γ_L , subject to $\gamma_L \leq 1/\gamma_U$.

First, assume we are given a fixed, arbitrary feasible set of bin boundaries $\{\beta_i\}$, a nonnegative value of κ_1^i for each bin, and that κ_2^i can take any real value. Because the problem is otherwise unconstrained, the constraint $\gamma_L \leq 1/\gamma_U$ will bind, or $\gamma_L = 1/\gamma_U$. Suppose γ_L is maximal with respect to the relevant constraints, and that bin i is the source of γ_L ; i.e.,

$$\gamma_L = \frac{\beta_{i-1}(\kappa_1^i - \kappa_2^i)}{1 + \kappa_2^i \beta_{i-1}}. \quad (35)$$

Simultaneously, κ_2^i must satisfy the following (the only other constraint on κ_2^i):

$$\frac{\beta_i(\kappa_1^i - \kappa_2^i)}{1 + \kappa_2^i \beta_i} \geq \gamma_U. \quad (36)$$

It is clear that γ_L is decreasing in κ_2^i ; if γ_L is indeed optimal, κ_2^i must be binding the constraint in (37). Thus, a single bin

⁴In [14], it is shown that assuming $\kappa_1^i \geq \kappa_2^i$ is without loss of generality.

generates both γ_L and γ_U , and since $\gamma_L = 1/\gamma_U$, we have that

$$\frac{\beta_{i-1}(\kappa_1^i - \kappa_2^i)}{1 + \kappa_2^i \beta_{i-1}} = \frac{\beta_i(\kappa_1^i - \kappa_2^i)}{1 + \kappa_2^i \beta_i}, \quad (37)$$

which implies that

$$\kappa_2^i = \frac{(\kappa_1^i)^2 \beta_{i-1} \beta_i - 1}{\beta_{i-1} + \beta_i + 2(\kappa_1^i) \beta_{i-1} \beta_i} \quad (38)$$

Now, if we re-introduce the constraint that $\kappa_2^i \geq 0$, we find that it may no longer be possible to satisfy (39), but that since γ_L is decreasing in some κ_2^i , simply saturating κ_2^i at 0 still yields a maximal γ_L while respecting $\gamma_L \leq 1/\gamma_U$. Thus, for any binning, choosing κ_2^i according to (27) is sufficient to ensure an optimal price of anarchy.

Finally, with (27), it is simple to show that γ_L is non-decreasing in each of the κ_1^i coefficients, so letting $\kappa_1^i = \kappa_{\max}$ suffices to minimize the price of anarchy. ■

Proof of Theorem 3: First, note that Optimization Problem (P) is largely a re-statement of Lemmas 3.1 and 3.2, with the sole exception of constraint (23). By the definition of γ_L in Lemma 3.1, it is clear why we desire to maximize γ_L . By the fact regarding γ_L in Lemma 3.2, it is clear that any optimal binning will satisfy constraint (22).

The only curious element of the optimization problem is constraint (23), which we show here does not reduce the optimality of the solutions, even when the constraint is active. This constraint plays a role in cases when κ_{\max} is small, and avoids the non-smoothness of κ_2^i in (27) and γ_L in (28). Let $\{\beta_i^*\}$ be an optimal solution to the optimization problem in which constraint (23) binds. Then $\beta_1^* = 1/(\kappa_{\max}^2 S_L)$, or $\kappa_{\max} = 1/\sqrt{S_L \beta_1^*}$. Note that at this point, according to (27), $\kappa_2^1 = 0$, and this is the precise breakpoint at which the expression for κ_2^1 “switches over” from 0 to the non-constant function of κ_{\max} . Thus, the first effect of constraint (23) is that it ensures that κ_2^1 will always be a smooth function of the bin boundaries.

Second, considering (28), note that $\beta_1 = 1/(\kappa_{\max}^2 S_L)$ is also the precise breakpoint at which γ_L “switches over” from $\beta_{i-1} \kappa_{\max}$ to $(\kappa_{\max} + 1/\beta_i)/(\kappa_{\max} + 1/\beta_{i-1})$. Thus, the second effect of constraint (23) is that it ensures that γ_L itself will always be a consistent function of the bin boundaries.

To see that including (23) does not reduce the optimality of solutions, note that if it were true that $\beta_1^* < 1/(\kappa_{\max}^2 S_L)$, it would also necessarily be true that the price of anarchy would simply be determined by $\gamma_L = S_L \kappa_{\max}$. Thus, β_1 has no impact on the price of anarchy until it reaches the $1/(\kappa_{\max}^2 S_L)$ threshold.

The essential uniqueness of the solution of (P) follows from the fact that the only way to increase γ_L is to raise the lower-boundary of some bin, which simultaneously raises the upper-boundary of the next-lower bin, serving to decrease γ_L . If (23) is active at an optimizer of (P), $\{\beta_i\}_{i=2}^m$ are not unique, but all optimizers yield the same price of anarchy.

To show the price of anarchy bound, first note that for a fixed m and S_L , the price of anarchy will always be increasing in S_U , since this represents increasing the sensitivity-

uncertainty of our population. Our arguments thus involve investigating the limiting price of anarchy as $S_U \rightarrow \infty$.

First, suppose that $\kappa_{\max} < 1/S_L$, so that for some binnings we activate constraint (23). When (23) is active, the price of anarchy is determined only by the lowest-index bin; in the language of Lemma 3.1, $\gamma_L = S_L \kappa_{\max}$. This can be considered a worst-case situation, so we include it in the price of anarchy expression (21) via the μ argument.

However, even when $\kappa_{\max} < 1/S_L$, it is possible that constraint (23) will not be active, so we must consider the case when constraint (22) is active. In this case, each bin has the same “width,” or for $i = (m - 1)$, it is true that

$$\gamma_L = \frac{\kappa_{\max} + 1/\beta_{m-1}}{\kappa_{\max} + 1/\beta_{m-2}} = \frac{\kappa_{\max} + 1/S_U}{\kappa_{\max} + 1/\beta_{m-1}}.$$

In general, closed-form expressions for bin boundaries resulting from this are quite complicated, but considering a high- S_U case can simplify things considerably:

$$\lim_{S_U \rightarrow \infty} \gamma_L = \frac{\kappa_{\max} + 1/\beta_{m-1}}{\kappa_{\max} + 1/\beta_{m-2}} = \frac{\kappa_{\max}}{\kappa_{\max} + 1/\beta_{m-1}}. \quad (39)$$

Since the above is true for any $m > 1$, we can use it to inductively deduce the structure of an optimal binning for any positive S_L and infinite S_U . First, letting $m = 2$, fixing β_1 implies the following unique value for S_L :

$$S_L = \kappa_{\max} / ((\kappa_{\max} + 1/\beta_1)^2 - \kappa_{\max}^2).$$

Inductively, it can be shown that for arbitrary m and fixed β_{m-1} , the unique implied value of S_L is given by

$$S_L = \kappa_{\max}^{m-1} / ((\kappa_{\max} + 1/\beta_{m-1})^m - \kappa_{\max}^m).$$

Solving this for β_{m-1} , we have that for any S_L , κ_{\max} , m , and infinite S_U ,

$$\beta_{m-1} = \left(\kappa_{\max} \left[\left(\frac{1 + \kappa_{\max} S_L}{\kappa_{\max} S_L} \right)^{1/m} - 1 \right] \right)^{-1}. \quad (40)$$

This β_{m-1} represents the highest bin boundary in an optimal binning for an infinite- S_U population. To compute the corresponding γ_L^∞ , we can simply substitute (30) into (29) and simplify:

$$\gamma_L^\infty = \left(\frac{\kappa_{\max} S_L}{1 + \kappa_{\max} S_L} \right)^{1/m}, \quad (41)$$

the source of λ in the theorem statement. Any finite S_U will yield a better price of anarchy than an infinite one, so for any game, $\gamma_L \geq \min\{\kappa_{\max} S_L, \gamma_L^\infty\}$, so by Lemma 3.1 the expression (21) is valid for any S_U . ■

V. CONCLUSIONS

The model of price-discrimination we introduce in this paper demonstrates that price-discrimination may be a powerful tool for influencing social behavior in systems that are not perfectly characterized. Of course, our model represents a simplified view of discriminatory pricing, and future research includes a study of the obstacles to a successful implementation. Among these is the issue of how a designer might obtain

a binning; even for a small number of bins, agents may have an incentive to mis-represent their sensitivities, since doing so could result in paying lower taxes. Another interesting (and potentially troubling) issue is that of fairness in price-discrimination; one aspect of the “optimal” discriminatory prices of Theorem 3 is that low-sensitivity agents’ absolute costs are much lower than those of high-sensitivity agents.

REFERENCES

- [1] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [2] R. Johari and J. N. Tsitsiklis, “Efficiency Loss in a Network Resource Allocation Game,” *Mathematics of Operations Research*, vol. 29, pp. 407–435, Aug. 2004.
- [3] H. Youn, M. Gastner, and H. Jeong, “Price of Anarchy in Transportation Networks: Efficiency and Optimality Control,” *Physical Review Letters*, vol. 101, p. 128701, Sept. 2008.
- [4] P. N. Brown and J. R. Marden, “The Robustness of Marginal-Cost Taxes in Affine Congestion Games,” (*Under Review*), 2015.
- [5] M. Gairing, “Covering games: Approximation through non-cooperation,” in *Lecture Notes in Computer Science*, vol. 5929 LNCS, pp. 184–195, 2009.
- [6] H.-L. Chen, T. Roughgarden, and G. Valiant, “Designing Network Protocols for Good Equilibria,” 2010.
- [7] J. R. Marden and A. Wierman, “Distributed Welfare Games,” *Operations Research*, vol. 61, pp. 155–168, 2013.
- [8] R. Gopalakrishnan, J. R. Marden, and A. Wierman, “Potential games are necessary to ensure pure nash equilibria in cost sharing games,” *Proceedings of the fourteenth ACM conference on Electronic commerce - EC '13*, pp. 563–564, 2013.
- [9] W. Sandholm, “Negative Externalities and Evolutionary Implementation,” *The Review of Economic Studies*, vol. 72, pp. 885–915, July 2005.
- [10] L. Fleischer, K. Jain, and M. Mahdian, “Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games,” *45th Annu. IEEE Symp. on Foundations of Computer Science*, pp. 277–285, 2004.
- [11] G. Karakostas and S. Kolliopoulos, “Edge pricing of multicommodity networks for heterogeneous selfish users,” *45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 268–276, 2004.
- [12] M. Beckman, C. McGuire, and C. B. Winsten, “Studies in the Economics of Transportation,” 1956.
- [13] W. Sandholm, “Evolutionary Implementation and Congestion Pricing,” *The Review of Economic Studies*, vol. 69, no. 3, pp. 667–689, 2002.
- [14] P. N. Brown and J. R. Marden, “Optimal Mechanisms for Robust Coordination in Congestion Games,” in *54th IEEE Conference on Decision and Control*, p. (To Appear), 2015.
- [15] R. B. Dial, “Network-Optimized Road Pricing: Part I: A Parable and a Model,” *Operations Research*, vol. 47, pp. 54–64, Jan. 1999.
- [16] H. Varian, “Monopoly Behavior,” in *Intermediate Microeconomics, a Modern Approach, 8th ed.*, pp. 462–484, 2010.
- [17] W. J. Baumol, “Ramsey Pricing,” *The New Palgrave*, vol. 4, pp. 49–51, 1987.
- [18] K. Tsakalozos, H. Killapi, E. Sitaridi, M. Roussopoulos, D. Paparas, and A. Delis, “Flexible use of cloud resources through profit maximization and price discrimination,” in *Proceedings - International Conference on Data Engineering*, pp. 75–86, 2011.
- [19] H. Shen and T. Basar, “Optimal Nonlinear Pricing for a Monopolistic Network Service Provider with Complete and Incomplete Information,” *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 6, pp. 1216–1223, 2007.
- [20] D. Bergemann and K. H. Schlag, “Pricing without Priors,” *Journal of the European Economic Association*, vol. 6, no. 2-3, pp. 560–569, 2008.
- [21] A. Mas-Colell, “On a Theorem of Schmeidler,” *Mathematical Economics*, vol. 13, pp. 201–206, 1984.
- [22] R. B. Wilson, *Nonlinear Pricing*. Oxford University Press, 1993.