

Static Pricing for a Network Service Provider*

Felipe Caro[†]
David Simchi-Levi[‡]

March 26, 2006

Abstract

This article studies the static pricing problem of a network service provider who has a fixed capacity and faces different types of customers (classes). We consider a single-bandwidth tree network, meaning that each class can have its own capacity constraint but it is assumed that all classes have the same resource requirements. The provider must decide a static price for each class. The customer types are characterized by their arrival process, with a price-dependant arrival rate, and the random time they remain in the system. The goal is to characterize the optimal static prices in order to maximize the provider's revenue. We report new structural findings and insights, illustrative numerical examples, and alternative proofs for some known results. This problem was originally thought for a company that sells phone cards and needs to set the price-per-minute for each destination.

Keywords: phone cards, Erlang loss system, product-form, stochastic knapsack, quasiconcavity.

1. Introduction

This paper considers the case of a network service provider who has a fixed capacity (bandwidth) and faces different types of customers (classes). Each type of customers can have its own capacity constraint and the provider must decide a static price (per unit of time) for each class. The customer types are characterized by their arrival process, with a price-dependant arrival rate, and the random time they remain in the system. All classes are assumed to have the same bandwidth requirement. This can be seen as the case when all classes demand the same service level which can be satisfied by allocating a constant amount of resource in the corresponding links. A customer is “blocked” (i.e. cannot enter the system) and lost if not enough capacity is available when he arrives. Therefore, in this paper the network service provider can be seen as a loss system with a tree topology, meaning that all classes share a common link but also use a “branch” that is their own link.

Though many real-life situations could fit in this framework, for example a dial-up Internet provider or a call center, this problem was originally thought for a company that sells phone cards, and that is the application we keep in mind throughout the paper. Then, we have that customer types represent call destinations, for example, the phone calls to Chile on a weekday (or even more precise, to Santiago the first Monday of June from noon to 4pm). Our objective is to characterize the optimal static prices in order to maximize the provider's revenue.

*Research supported in part by the Center of eBusiness at MIT, ONR Contracts N00014-95-1-0232 and N00014-01-1-0146, and by NSF Contracts DMI-9732795, DMI-0085683 and DMI-0245352.

[†]UCLA Anderson School of Management, 110 Westwood Plaza Suite B-420, Los Angeles, CA 90095, USA. email: fcaro@anderson.ucla.edu

[‡]Dept. of Civil and Environmental Engineering and the Engineering Systems Division, MIT, 77 Massachusetts Ave. Rm. 1-171 Cambridge, MA 02139, USA. email: dslevi@mit.edu

We note that, given the finite capacity of the system, it is intuitive to believe that a dynamic pricing scheme (i.e. to charge a higher price when the system is heavily loaded) would substantially improve the provider's revenue. Nevertheless, besides implementation issues, instantaneous changes in prices are not well received among customers and can cause a loss in image to the company. Furthermore, there is theoretical evidence showing that static prices are asymptotically optimal for a specific time-window (see Paschalidis and Tsitsiklis [13] and Paschalidis and Liu [14]). Therefore, if the small extra gain obtained by using dynamic pricing does not offset the image costs, then finding the best static prices is the correct goal.

In terms of contributions, opposed to most of the literature that focuses on asymptotic regimes (cf. §2), this paper provides an intuitive characterization of the optimal static prices, we discuss how easily they can be found, and enrich the limited existing literature with new structural findings for particular, yet important, loss systems. Additionally, alternative proofs are given for some known results, and in general, further insight is provided for the pricing problem of a network service provider that helps identify when an exact or an asymptotic approach is more desirable.

The remainder of the paper has the following structure: The next section (§2) provides a brief literature review. Then, in §3, the special tree topology here considered is introduced and the static pricing model is formulated. All the theoretical results are derived in §4. The asymptotically optimal prices are introduced in §5. Some illustrative numerical examples are given in §6, and finally §7 summarizes the conclusions and possible extensions to more complex settings.

2. Literature Review

There is a vast literature on loss systems, and most of the effort has concentrated in finding structural properties of the blocking probabilities, answering design questions and/or focusing on admission control. Two important references that summarize large streams of previous works are Kelly [8] and Ross [15]. This paper benefits from many of the results found in there.

Loss systems are closely related to queueing models, where opposed to our case, pricing has been studied assuming infinite buffers. The absence of blocking brings congestion and most of the work is interested in relating prices with customer delay (waiting times). Mendelson [11] is one of the first papers that specifically looks at this issue in a simple queueing model for a single customer class. Since then many extensions to a multi-class setting have been reported. For instance, see Ha [6] and the references therein.

Pricing in loss systems has received much less attention. Courcoubetis and Reiman [4] propose models for a single link with N circuits and consider substitution effects between different types of calls. Results are given for an asymptotic regime where both, the available capacity and potential load, grow to infinity. Another attempt is given by Lanning *et al* [9], where a dynamic pricing model is developed for an Erlang blocking system (i.e. single link and single customer class). The arrival rate depends not only on the price but also on its percentage change. The analysis is confined to the infinite capacity case. Ziya *et al* [19] present a static pricing problem for a service facility that is modeled as a single-class single-server queue that might have limited waiting room (hence, customers are blocked). The results establish the relation of the optimal price with respect to the system parameters. The model introduced by Carrizosa *et al* [3] also deserves attention. Despite the fact that their formulation is a probabilistic admission control problem, there are some

similarities with the results given in this paper and will be noted further on.

For our purposes, Paschalidis and Tsitsiklis [13] provide the most relevant reference in the previous stream. Their dynamic pricing model is also for a single link but with multiple customer types that might require different units of bandwidth. As mentioned before, their results show that static prices are asymptotically optimal. The definition of asymptotic is for the case of *many relatively small users*, indicating the regime when the number of users grows to infinity but so does the capacity. The extension of the asymptotic optimality to the network setting, which would fit our case, is almost direct and is done by Paschalidis and Liu [14]. They also provide some notions on the rate of convergence, and show that the asymptotic optimality still holds when substitution is allowed as in Courcoubetis and Reiman [4].

From another perspective, the literature on pricing broadband IP networks provides several other pricing schemes that can be seen as alternatives to the static one presented in this paper. This includes, among others, smart-market pricing, effective bandwidth pricing, and proportional fairness pricing. We refer the interested reader to the book by Courcoubetis and Weber [5].

3. The Model

The system modeled in this paper is represented in Figure 1. Let $k = 1, \dots, K$ represent the different call destinations. In what follows, by class- k customers we refer to those that make calls to destination k . The network has a common link with N lines (circuits) that receives all incoming service requests. Then, for each destination k there is an outbound link with N_k lines. Let $\mathbf{C} = (N, N_1, \dots, N_K)$ denote the capacity vector (note that throughout the paper boldface symbols represent vectors).

In our motivating example of a company that sells phone cards, the common link represents the physical switch where all incoming calls are received, and the capacity N corresponds to the maximum number of simultaneous calls that the switch can handle. The outbound links represent the different carriers that complete the corresponding calls and the capacity N_k comes from the contract signed with a carrier for a specific destination. Our model takes the capacity vector \mathbf{C} as an input. In other words, we assume that the network was designed at a tactical level and the only operational decision is the vector of prices.

PLACE FIGURE 1 HERE

It is natural to assume that all destinations have the same resource requirements, so each call to destination k uses one line of the common link and one line of the k -th outbound link. Then what we have is a tree network. It is clear that the model remains the same if the class-specific links are located before the common link as in an access network (see [15]).

Let p_k be the price per unit of time charged to class- k customers. Class- k customers arrive to the system according to a Poisson process with rate $\lambda_k(p_k)$ and their holding time is exponential with mean $1/\mu_k$. We note that the given price dependence has three implicit assumptions:

1. *The arrival rate of class k depends only on its own price.* Therefore, opposed to Courcoubetis and Reiman [4], we do not consider substitution effects among classes. Again, this is reasonable in the case of a phone card company since *variations in the price of calls to China does*

not affect the number of calls made to Chile. The relaxation of this assumption is discussed at the end of the paper.

2. *We are assuming that each class acts as a single decision maker.* In other words, all the members of a class set the aggregate arrival rate collusively. This is a common assumption when the goal is to maximize total revenue. An alternative approach to ours is an “atomistic” model (see [12]) where every potential user decides individually if he/she joins the system. Then the total arrival rate the service provider sees is the aggregation of many individual decisions.
3. *The holding times are price-independent.* This assumption (which is also made in [9], [11], [12], and [13]) is required for analysis tractability (so that arrivals and service times are independent) and can be considered one of the limitations of our model.

For each destination k we let the price vary within a closed interval $[0, p_k^{max}]$ such that $\lambda_k(p_k^{max})$ is arbitrarily close to 0. The functions $\lambda_k(p_k)$ are assumed to be continuously differentiable and strictly decreasing with inverse $p_k(\lambda_k)$. Let $y_k(p_k) = \lambda_k(p_k)/\mu_k$ be the traffic intensity of class k , and $\mathbf{y}(\mathbf{p})$ is the corresponding vector with all customer classes. We assume that for every vector of prices \mathbf{p} the arrival processes of different classes are independent of each other.

We model the tree topology of Figure 1 as a loss network (see [15]). Let \mathbf{n} be the state vector of size K where the k -th component n_k represents the number of class- k calls in the system. The natural admission policy is: admit a class- k call if there is one line available in the common link and one line available in the k -th outbound link. Then the state space \mathbf{S} is given by all the nonnegative integer vectors \mathbf{n} such that $n_k \leq N_k \forall k = 1, \dots, K$ and $\mathbf{n}'\mathbf{e} \leq N$, where \mathbf{e} is a vector of ones. In more general terms, this admission policy is equivalent to a loss system operating under a *threshold policy* and is part of a greater family known as *the coordinate convex policies*¹. In turn, the set of *threshold policies* contains two particular cases that will be of further interest: (i) *the complete sharing policy*: corresponds to the case when $N_k \geq N, \forall k$; and (ii) *the complete partitioning policy*: corresponds to the case when $N_1 + N_2 + \dots + N_K \leq N$ so that the tree network decouples into K single-class loss systems.

Let X_k denote the random number of class- k calls in the system in steady state. Since there is a one-to-one mapping between traffic intensities \mathbf{y} and static prices \mathbf{p} , depending on the context we will write $X_k(\mathbf{y})$ or $X_k(\mathbf{p})$. In the same sense, when necessary we will explicitly write the dependence on the capacity vector \mathbf{C} . Using standard reversibility arguments (see [8] or [15]) it can be verified that the equilibrium distribution for the loss network (given \mathbf{p}) is:

$$\Pr(\mathbf{X}(\mathbf{p}) = \mathbf{n}) = \frac{1}{G(\mathbf{p})} \prod_{k=1}^K \frac{y_k(p_k)^{n_k}}{n_k!}, \quad \mathbf{n} \in \mathbf{S}, \quad \text{where } G(\mathbf{p}) = \sum_{\mathbf{n} \in \mathbf{S}} \prod_{k=1}^K \frac{y_k(p_k)^{n_k}}{n_k!}. \quad (1)$$

It is well known that the above product form distribution remains the same if instead of exponentials the holding times have general distributions with finite means $1/\mu_k$, for each class k . Also, from the continuity properties of $\mathbf{y}(\mathbf{p})$ we have that $\Pr(\mathbf{X}(\mathbf{p}))$ is continuously differentiable with respect to \mathbf{p} (or \mathbf{y}) and so is $E[\mathbf{X}(\mathbf{p})]$.

¹Roughly speaking, in our context, an admission control policy is *coordinate convex* if the equilibrium probabilities have a normalized product form. See [15] for a formal definition.

An incoming call that does not find enough capacity available in the system will be rejected (blocked) and leaves without being served (the standard renewal assumption for the arrival process ignores retrials). Let $B_k(\mathbf{p})$ and $\bar{B}_k(\mathbf{p})$ be the class- k blocking and nonblocking probabilities respectively. These quantities can be calculated numerically via efficient algorithms for a given vector of prices (see, for example, see Tables 1-4 in [16] where strong polynomial complexity bounds and numerical results are given for several exact procedures). We note and emphasize that this a particular characteristic of the special tree topology we consider. Computing blocking probabilities in arbitrary topologies is an NP-complete problem (see [15]).

With no loss of generality, we assume that the cost per unit of time of a class- k call is equal to zero. Then, given the capacity vector \mathbf{C} , the steady state revenue function that the service provider wants to maximize over all $\mathbf{p} \in \mathbb{R}_+$ can be expressed as:

$$W(\mathbf{p}; \mathbf{C}) = \sum_{k=1}^K p_k E[X_k(\mathbf{p})] = \sum_{k=1}^K p_k y_k(p_k) \bar{B}_k(\mathbf{p}) \quad (2)$$

The last equality follows from Little's Law and shows that charging class k a price p_k per unit of time is the same as charging them a flat fee p_k/μ_k upon arrival

The analysis is simplified if instead of setting prices, the service provider decides the arrival rate λ_k for each class and then charges the price dictated by the inverse of the demand function $\lambda_k(p_k)$. We denote such inverse function as $p_k(\lambda_k)$ defined on the closed interval $[0, \lambda_k^{max}]$, where $\lambda_k^{max} = \lambda_k(0)$.² Let $R_k(\lambda_k) \equiv \lambda_k \cdot p_k(\lambda_k)$ be the instantaneous reward rate. We assume that $R'_k(0) > 0 \forall k$ to exclude trivial solutions.

Let $\Lambda = [0, \lambda_1^{max}] \times \dots \times [0, \lambda_K^{max}]$. Then, the goal is to solve and characterize the solution of the following optimization problem:

$$(P) \quad \max_{\boldsymbol{\lambda} \in \Lambda} W(\boldsymbol{\lambda}; \mathbf{C}) = \frac{1}{G(\boldsymbol{\lambda}; \mathbf{C})} \sum_{k=1}^K R_k(\lambda_k) \cdot \frac{\partial G(\boldsymbol{\lambda}; \mathbf{C})}{\partial \lambda_k}. \quad (3)$$

Note that we are using the identities $\frac{\partial G(\mathbf{y}; \mathbf{C})}{\partial y_i} = G(\mathbf{y}; \mathbf{C} - \mathbf{b}_i)$ and $\bar{B}_i(\mathbf{y}) = G(\mathbf{y}; \mathbf{C} - \mathbf{b}_i)/G(\mathbf{y}; \mathbf{C})$, where \mathbf{b}_i corresponds to a vector with $(K+1)$ components, all equal to zero except for the first and the $(i+1)$ -th position which are equal to one.

A formulation equivalent to (P) can be obtained if we define $\lambda_k = x_k \cdot \lambda_k^{max}$ and consider $x_k \in [0, 1] \forall k$ as the decision variables. This is how Carrizosa *et al* [3] proceed. They regard their problem as finding, for each class, the rejection probability $(1 - x_k)$ when there is an idle server.

Finally, it is important to note that with an appropriate definition of the $R_k(\lambda_k)$ functions, our model could also fit a social welfare maximization problem. One example is the *consumer surplus* model proposed by Paschalidis and Tsitsiklis [13].

4. Characterization of the Optimal Static Prices

4.1 First-order conditions. The existence of optimal static prices, or equivalently, of optimal static arrival rates, is guaranteed by the fact that we are maximizing a continuous functions over the compact set Λ . To begin with, we will provide first-order conditions for the optimal static prices, and then, we will analyze second-order properties. Directly differentiating the revenue function in equation (3) with respect to λ_i , and assuming that $N_i > 0$ so that $\frac{\partial G}{\partial \lambda_i} \neq 0$, we obtain:

²If $\lambda_k(0) = \infty$, then λ_k^{max} is sufficiently large.

$$\begin{aligned}
\frac{\partial W(\boldsymbol{\lambda}; \mathbf{C})}{\partial \lambda_i} &= \frac{1}{G(\boldsymbol{\lambda}; \mathbf{C})^2} \left(\left(R'_i(\lambda_i) \frac{\partial G}{\partial \lambda_i} + \sum_{k=1}^K R_k(\lambda_k) \frac{\partial^2 G}{\partial \lambda_i \partial \lambda_k} \right) G(\boldsymbol{\lambda}; \mathbf{C}) - \left(\sum_{k=1}^K R_k(\lambda_k) \frac{\partial G}{\partial \lambda_k} \right) \frac{\partial G}{\partial \lambda_i} \right) \\
&= \frac{\frac{\partial G}{\partial \lambda_i}}{G(\boldsymbol{\lambda}; \mathbf{C})} \left[R'_i(\lambda_i) - \left(\sum_{k=1}^K R_k(\lambda_k) \frac{\frac{\partial G}{\partial \lambda_k}}{G(\boldsymbol{\lambda}; \mathbf{C})} - \sum_{k=1}^K R_k(\lambda_k) \frac{\frac{\partial^2 G}{\partial \lambda_k \partial \lambda_i}}{\frac{\partial G}{\partial \lambda_i}} \right) \right] \\
&= \frac{\bar{B}_i(\boldsymbol{\lambda}; \mathbf{C})}{\mu_i} [R'_i(\lambda_i) - (W(\boldsymbol{\lambda}; \mathbf{C}) - W(\boldsymbol{\lambda}; \mathbf{C} - \mathbf{b}_i))]
\end{aligned} \tag{4}$$

where in the second equation we have used the equality of the second-order cross-derivatives of $G(\boldsymbol{\lambda}; \mathbf{C})$, and the last equation follows from the identity $\partial G / \partial \lambda_i = \partial G / \partial y_i \cdot \mu_i^{-1}$. Then, from the formula above, we have a direct result that gives the first-order conditions for the optimal static arrival rates (and consequently for the optimal static prices).

Proposition 1 *Let λ^* be a vector of static arrival rates that solves the maximization problem (P), then it must satisfy the following first-order conditions:³*

$$\begin{aligned}
\text{If } \lambda_i^* > 0 \text{ then } R'_i(\lambda_i^*) &= W(\boldsymbol{\lambda}^*; \mathbf{C}) - W(\boldsymbol{\lambda}^*; \mathbf{C} - \mathbf{b}_i) \\
\text{If } \lambda_i^* = 0 \text{ then } R'_i(0) &\leq W(\boldsymbol{\lambda}^*; \mathbf{C}) - W(\boldsymbol{\lambda}^*; \mathbf{C} - \mathbf{b}_i)
\end{aligned} \tag{5}$$

The conditions given by (5) say that at the optimum, if class i is active (i.e. $\lambda_i^* > 0$), then the marginal revenue of class i , $R'_i(\lambda_i^*)$, must be equal to the opportunity cost of admitting a call of class i and therefore blocking future calls of other classes. In steady-state this opportunity cost is given by the difference between the revenue of the original system, $W(\boldsymbol{\lambda}^*; \mathbf{C})$, and the revenue of another system that is identical to the original but with one line less in the common link and one line less in the i -th outbound link (i.e. a system with capacities $(\mathbf{C} - \mathbf{b}_i)$).⁴

PLACE FIGURE 2 HERE

The derivative formula (4) is useful for implementing greedy optimization routines. Unfortunately, in general the revenue function is not concave. In Figure 2 we plot an example for two classes with linear demands. Depending on the parameters the function looks perfectly concave (on the left side) or clearly non concave (on the right side). Nevertheless, as we will shortly prove, the latter can still be quasiconcave.

4.2 Single-class and complete partitioning policy. We start the second-order analysis with a single customer class. For that case, the system reduces to an $M/G/N/N$ queue, also known as Erlang's loss system. We first give a useful Lemma that is well known in the literature (for instance, see [3]). However, we provide a simpler proof that is interesting per se.

Lemma 1 *In an $M/G/N/N$ system, the inverse of the nonblocking probability is convex with respect to the traffic intensity.*

³We exclude the case $\lambda_i^* = \lambda_i^{max}$ since this corresponds to a zero price which will never be optimal.

⁴Formulas similar to (4) can be found in [15].

Proof: For $N = 1$ the function is linear, and the result is trivial. Assume $N > 1$. It can be verified directly that the second derivative being non-negative is equivalent to having:

$$\left[(G(y; N-1)G(y; N-2) + G(y; N)G(y; N-3))G(y; N-1)^2 - 2 \cdot G(y; N)G(y; N-1)G(y; N-2)^2 \right] \leq 0$$

Dividing the above expression by $G(y; N)G(y; N-1)^2G(y; N-2)$ and rearranging terms we see that convexity will hold if and only if the following inequality holds:

$$\underbrace{\frac{G(y; N-1)}{G(y; N)}}_{\bar{B}(y; N)} + \underbrace{\frac{G(y; N-3)}{G(y; N-2)}}_{\bar{B}(y; N-2)} \leq 2 \cdot \underbrace{\frac{G(y; N-2)}{G(y; N-1)}}_{\bar{B}(y; N-1)}$$

Then the result follows from Ross and Yao [17] since the above inequality is true due to the concavity of the nonblocking probabilities with respect to N (recall that $N > 1$). \square

An remarkable fact comes as a by-product in the proof of Lemma 1. We have shown that, for an $M/G/N/N$ system, the (discrete) concavity of the nonblocking probability with respect to the capacity of the common link N is equivalent to the (continuous) convexity of its reciprocal with respect to the traffic intensity y . We now use this result in our context.

Proposition 2 *For the single-class case:*

- *If the inverse demand function $p(\lambda)$ is (strictly) concave, then the revenue function is (strictly) concave.*
- *If the instantaneous reward function $R(\lambda)$ is (strictly) concave, then the revenue function is (strictly) quasiconcave.⁵*

Proof: For the single-class case revenue function simplifies to $W(\lambda) = p(\lambda) \cdot E[X(\lambda)] = R(\lambda) \cdot \bar{B}(\lambda)/\mu$. For ease of notation we omit the capacity N .

When $p(\lambda)$ is (strictly) concave, the result follows from the fact that we use the fact that $W(\lambda)$ is equal to the product of two positive functions, one (strictly) concave and decreasing, and the other also concave but non-decreasing.⁶

For the second case, we use Lemma 1. The desired quasiconcavity result follows since the revenue function $W(\lambda)$ can be written as the ratio of $R(\lambda)/\mu$, that is non-negative (strictly) concave, and $\bar{B}(\lambda)^{-1}$, that is positive convex (see [2]). \square

Under the assumptions of Proposition 2, the revenue function is also pseudoconcave, which is enough for the first-order conditions (5) to become sufficient (see section 3.6 of [2]). When strict quasiconcavity holds, problem (P) is guaranteed to have a unique optimal solution.

The result given above for the single-class case can be directly extended to the case with multiple classes when the revenue function is separable, or equivalently, when the loss system operates under a *complete partitioning policy*:

⁵From Avriel *et al* [2], a real function $f(x)$ is strictly quasiconcave on a convex set S if $\forall x_1, x_2 \in S$, with $x_1 \neq x_2$, $f(\alpha x_1 + (1 - \alpha)x_2) > \min\{f(x_1), f(x_2)\} \forall \alpha \in (0, 1)$.

⁶The fact that $E[X(\lambda)]$ is non-decreasing follows from Proposition 1 in [17]. For concavity, see [7].

Corollary 1 For a multiple-class loss system operating under a complete partitioning policy:

- If $p_k(\lambda_k)$ is (strictly) concave $\forall k$, then the revenue function is (strictly) concave.
- If $R_k(\lambda_k)$ is (strictly) concave $\forall k$, then the revenue function is (strictly) quasiconcave.

Proof: The first case is trivial. The case with all $R_k(\lambda_k)$ concave is not obvious since the sum of quasiconcave functions in general is not quasiconcave, but an exception is given in Almogly and Levin [1]. See the Corollary of Theorem 5 which can be applied here because $\bar{B}(\lambda_k, N_k)^{-1} \geq 1$. \square

When all demands are linear, the first case in Corollary 1 follows. This is an important particular case but it could be too restrictive. The requirement that the instantaneous rewards $R_k(\lambda_k)$ are concave is much more general. For example it includes the case when $p_k(\lambda_k) = a_k \cdot \lambda_k^{-b_k}$ with $a_k > 0$ and $0 < b_k < 1$, or when $\lambda_k(p_k) = a_k \cdot \exp(-b_k p_k)$ with $a_k, b_k > 0$.

4.3 Complete sharing policy. We will now prove that when the loss system operates under a *complete sharing policy* the quasiconcavity of the revenue function remains. In this case the blocking probability is the same for all classes due to only sharing the common link. In fact, the system is equivalent to an $M/G/N/N$ queue with traffic intensity equal to $\sum_{k=1}^K \lambda_k / \mu_k$.

Proposition 3 Under complete sharing, if the instantaneous reward functions $R_k(\lambda_k)$ are (strictly) concave $\forall k$, then the revenue function is (strictly) quasiconcave.

Proof: For the *complete sharing* case all nonblocking probabilities are equal and the revenue function can be written as: $W(\boldsymbol{\lambda}; N) = (\sum_{k=1}^K \frac{R_k(\lambda_k)}{\mu_k}) / (G(\boldsymbol{\lambda}; N) / G(\boldsymbol{\lambda}; N - 1))$

As in the proof of Proposition 2, the only step that must be justified is the convexity of the inverse of the nonblocking probability $G(\boldsymbol{\lambda}; N) / G(\boldsymbol{\lambda}; N - 1)$, which follows from Lemma 1 and the composition of a convex function and the linear equation $y = \sum_{k=1}^K \lambda_k / \mu_k$. \square

Under a *complete sharing policy* the first order conditions (5) dictate that, at the optimum, all active classes must yield the same marginal revenue (equal to β^*). Clearly, it must be that $\beta^* > 0$. These remarks allow an important simplification of problem (P). For the rest of this subsection we assume that the instantaneous revenue functions $R_k(\lambda_k)$ are all strictly concave. As a consequence the optimal solution is unique. Consider the following parametrization of the arrival rates:

$$\lambda_k^P(\beta) = \begin{cases} R_k'^{-1}(\beta) & \text{if } \beta < R_k'(0) \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Let $\beta_k = R_k'(0)$. With no loss of generality, we assume that $\beta_1 \geq \beta_2 \geq \dots \geq \beta_K \geq \beta_{K+1} \equiv 0$. It can be shown that if $\boldsymbol{\lambda}^*$ is an optimal solution of (P), then there exists some $\beta^* \in [0, \beta_1]$ such that $\boldsymbol{\lambda}^P(\beta^*) = \boldsymbol{\lambda}^*$. We now provide a series of basic lemmas that show how problem (P) can be simplified.

Consider the one-dimensional problem (P_1) defined as the maximization of the parametric revenue function $W(\boldsymbol{\lambda}^P(\beta))$ subject to β in the interval $[0, \beta_1]$. The following Lemma is a direct consequence of the parametrization (6).

Lemma 2 *Problems (P) and (P₁) are equivalent. Moreover, the solution to (P₁) is unique.*

Lemma 2 shows that for the *complete sharing* case finding static prices can be reduced to a one-dimensional problem. A similar result is found in [3]. In our setting their objective function is equivalent to having $R_k(\lambda_k) = \lambda_k \forall k$. Their conclusions depend on this specific linear form, which does not fit the framework given by the parametrization (6). Therefore, despite similarities, neither result implies the other.

In what remains of this subsection we will derive some properties of problem (P₁). It is easy to verify that the parametric revenue function $W(\boldsymbol{\lambda}^P(\beta))$ is continuous on $[0, \beta_1]$ but in general is not differentiable at the break points β_k . However, the left and right derivatives always exist for any interior point. The top graph in Figure 3 shows the shape of $W(\boldsymbol{\lambda}^P(\beta))$ for different values of N and linear demands such that $\beta_1 = 30$ and $\beta_2 = 8$ (the latter is marked with a solid vertical line and corresponds to a non-differentiable point).

PLACE FIGURE 3 HERE

Let $\Delta W(\boldsymbol{\lambda}^P(\beta); N) \equiv W(\boldsymbol{\lambda}^P(\beta); N) - W(\boldsymbol{\lambda}^P(\beta); N-1)$ be the backward difference of $W(\boldsymbol{\lambda}^P(\beta); N)$ with respect to the capacity N . The next Lemma shows that finding the optimal solution of (P₁) is equivalent to finding a fixed point of $\Delta W(\boldsymbol{\lambda}^P(\beta))$.

Lemma 3 *β^* is optimal for (P₁) if and only if $\beta^* = \Delta W(\boldsymbol{\lambda}^P(\beta^*))$.*

Proof: First, suppose β^* is optimal for (P₁). Then, from Lemma 2, $\boldsymbol{\lambda}^P(\beta^*)$ is optimal for (P). Since $\boldsymbol{\lambda}^P(\beta^*) \neq 0$, there exists class i such that $\lambda_i^P(\beta^*) > 0$. The first-order conditions (5) imply that $R'_i(\lambda_i^P(\beta^*)) = \Delta W(\boldsymbol{\lambda}^P(\beta^*))$. Using (6) we obtain the desired equality.

For the converse, on the one hand, suppose that $\lambda_i^P(\beta^*) > 0$. Then $\lambda_i^P(\beta^*) = R_i'^{-1}(\beta^*)$ (from (6)), and consequently $R'_i(\lambda_i^P(\beta^*)) = \beta^* = \Delta W(\boldsymbol{\lambda}^P(\beta^*))$. On the other hand, suppose that $\lambda_i^P(\beta^*) = 0$. Then $\beta^* \geq R'_i(0)$ (from (6)), which means that $\Delta W(\boldsymbol{\lambda}^P(\beta^*)) \geq R'_i(0)$. Therefore, $\boldsymbol{\lambda}^P(\beta^*)$ satisfies the first-order conditions for (P). Since the proof of Proposition 3 also shows that (P) has a pseudoconcave objective function, it means that $\boldsymbol{\lambda}^P(\beta^*)$ is optimal for (P). Then, from Lemma 2, we have that β^* is optimal for (P₁), which completes the proof. \square

A few remarks with respect to Lemma 3. First, the function $\Delta W(\boldsymbol{\lambda}^P(\beta))$ is not necessarily contractive. However, Lemma 2 guarantees the existence of a unique fixed point. Second, from Lemma 3 it can be concluded that the optimal solution of (P₁) is an interior point. In fact, let $\lambda_k^r = \operatorname{argmax}_{\lambda_k} R_k(\lambda_k)$, then $\Delta W(\boldsymbol{\lambda}^P(0)) = (\sum_{k=1}^K R_k(\lambda_k^r)/\mu_k) \cdot (\bar{B}(\boldsymbol{\lambda}^r; N) - \bar{B}(\boldsymbol{\lambda}^r; N-1)) > 0$, and $\Delta W(\boldsymbol{\lambda}^P(\beta_1)) = 0$ (recall $\beta_1 > 0$). Consequently, neither extremes of the interval $[0, \beta_1]$ can be the solution of (P₁).

In the bottom graph of Figure 3 we plot $\Delta W(\boldsymbol{\lambda}^P(\beta))$ for the same instances than the top graph. As stated by Lemma 3, the optimal value of $W(\boldsymbol{\lambda}^P(\beta))$ is obtained when $\Delta W(\boldsymbol{\lambda}^P(\beta))$ intersects the 45° line.⁷ If β^* is the unique fixed point, then from the remark above we have that $\beta^* < \beta_1$, and therefore class 1 is always active, i.e. its optimal arrival rate is nonzero. If $\beta^* < \beta_2$, then class

⁷For expositional purposes, the axes of Figure 3 are not to scale.

2 is active, and so on. In general, all classes “to the right” of β^* will be active. For example, in Figure 3, class 2 has a zero arrival rate when $N = 10$ but becomes active when $N = 15$.

The top graph of Figure 3 suggests that, despite the non-differentiable points, the parametric revenue function is still quasiconcave. This fact is verified by our next Lemma.

Lemma 4 *The parametric revenue function $W(\boldsymbol{\lambda}^P(\beta))$ is quasiconcave.*

Proof: We will prove that $W(\boldsymbol{\lambda}^P(\beta))$ is unimodal, which in this case is equivalent to quasiconcavity (see Proposition 3.8 of [2]). Consider $\hat{\beta} \in (0, \beta_1)$. Through direct calculation it can be verified that the left derivative of $W(\boldsymbol{\lambda}^P(\beta))$ at $\hat{\beta}$ is given by:

$$\frac{\partial W(\boldsymbol{\lambda}^P(\hat{\beta}^-))}{\partial \beta} = \bar{B}(\boldsymbol{\lambda}^P(\hat{\beta})) \cdot (\hat{\beta} - \Delta W(\boldsymbol{\lambda}^P(\hat{\beta}))) \cdot \left[\sum_{i=1}^K \frac{1}{\mu_i} \frac{\partial \lambda_i^P(\hat{\beta}^-)}{\partial \beta} \right]. \quad (7)$$

The right derivative is the same but replacing $\hat{\beta}^-$ with $\hat{\beta}^+$. Because of the strict concavity of all R_i and the parametric definition (6), we have that for any class i :

$$\frac{\partial \lambda_i^P(\hat{\beta}^-)}{\partial \beta} \leq \frac{\partial \lambda_i^P(\hat{\beta}^+)}{\partial \beta} \leq 0 \quad \forall \hat{\beta} \in (0, \beta_1) \quad (8)$$

where the last inequality is strict if class i is active at $\hat{\beta}$. Since $\Delta W(\boldsymbol{\lambda}^P(0)) > 0$ and $\Delta W(\boldsymbol{\lambda}^P(\beta))$ is continuous with β^* its unique fixed point (cf. Lemmas 2 and 3), when $\hat{\beta} < \beta^*$ we have that $\hat{\beta} < \Delta W(\boldsymbol{\lambda}^P(\hat{\beta}))$, and then from equation (7):

$$\frac{\partial W(\boldsymbol{\lambda}^P(\hat{\beta}^-))}{\partial \beta} \geq \frac{\partial W(\boldsymbol{\lambda}^P(\hat{\beta}^+))}{\partial \beta} > 0. \quad (9)$$

If $\hat{\beta} > \beta^*$ then the inequalities in (9) are reversed. Summarizing, “before” β^* the function $W(\boldsymbol{\lambda}^P(\beta))$ is increasing, and “after” β^* it is decreasing. Hence, $W(\boldsymbol{\lambda}^P(\beta))$ is unimodal. \square

Lemma 4 is important since it states that the pricing problem for the *complete sharing* case can be solved using standard line search algorithms.

We conclude this subsection proving the monotonicity of the static prices with respect to N . It will be shown that this is a direct consequence of the concavity of the nonblocking probabilities with respect to that parameter. Let $p_k^*(N)$ denote the optimal static price of class k when the capacity of the common link is N . Then we have the following Proposition.

Proposition 4 *Consider a multi-class loss system operating under a complete sharing policy such that the instantaneous reward functions R_k are all strictly concave. Let N be a positive integer, then $p_k^*(N+1) \leq p_k^*(N) \forall k$.*

Proof: We will show that the optimal solution of (P_1) is non-increasing with respect to N , which implies the desired result.

Let β_N and β_{N+1} be the optimal solution of (P_1) for N and $N+1$ respectively. The concavity of the nonblocking probabilities implies that:

$$\Delta W(\boldsymbol{\lambda}^P(\beta); N) \geq \Delta W(\boldsymbol{\lambda}^P(\beta); N+1) \quad \forall \beta \in [0, \beta_1] \quad (10)$$

In other words, $\Delta W(\boldsymbol{\lambda}^P(\beta); N)$ “lies above” $\Delta W(\boldsymbol{\lambda}^P(\beta); N + 1)$. This remark, together with the fact that $\Delta W(\boldsymbol{\lambda}^P(0); N) > 0$, are enough for the intuition of why $\beta_N \geq \beta_{N+1}$ must hold (see for example Figure 3). However, we continue with a formal proof in order to show that the concavity of the nonblocking probability has even stronger consequences. In fact, suppose that $\beta_N < \beta_{N+1}$. From Lemma 3 the following right derivative formula at β_N can be verified for $N > 1$:

$$\frac{\partial \Delta W(\boldsymbol{\lambda}^P(\beta_N^+); N)}{\partial \beta} = -\bar{B}(\boldsymbol{\lambda}^P(\beta_N)) \cdot \left(\Delta W(\boldsymbol{\lambda}^P(\beta_N); N) - \Delta W(\boldsymbol{\lambda}^P(\beta_N); N-1) \right) \cdot \left[\sum_{i=1}^K \frac{1}{\mu_i} \frac{\partial \lambda_i^P(\beta_N^+)}{\partial \beta} \right]. \quad (11)$$

For $N = 1$ the derivative is equal to zero since $\Delta W(\boldsymbol{\lambda}^P(\beta); 1) = W(\boldsymbol{\lambda}^P(\beta))$. The concavity of the nonblocking probability (10) implies that the right derivative (11) is non-positive (actually negative for $N > 1$). Then, there must exist $\hat{\beta} > \beta_N$ such that $\Delta W(\boldsymbol{\lambda}^P(\hat{\beta}); N) < \hat{\beta}$. Let $q(\beta) \equiv \Delta W(\boldsymbol{\lambda}^P(\beta); N) - \beta$, which is a continuous function of β . On the one hand we have $q(\hat{\beta}) < 0$, and on the other hand, from Lemma 3 and (10), we have that $q(\beta_{N+1}) \geq 0$. Then there must exist $\xi \in [\min(\hat{\beta}, \beta_{N+1}), \max(\hat{\beta}, \beta_{N+1})]$ such that $q(\xi) = 0$. This implies $\xi = \beta_N$ (from the uniqueness of β_N), which is a contradiction since $\xi \geq \min(\hat{\beta}, \beta_{N+1}) > \beta_N$. Consequently, it must be that $\beta_N \geq \beta_{N+1}$, and the proof is complete. \square

4.4 Threshold policy.

A *threshold policy* corresponds to the most general case for the network service provider depicted in Figure 1. To exclude the particular cases covered in the previous subsections, we assume that $N_k \leq N \forall k$ and $\sum_{k=1}^K N_k > N$. We would like to identify when the following two conditions hold: (i) the level sets of the revenue function are connected, and (ii) the optimal solution of (P) is unique. Condition (i) is necessary (and sufficient under proper assumptions) for greedy search algorithms to converge to global maxima. Condition (ii) is a desirable property that allows the construction of one-to-one mappings as we did in § 4.3. It can be shown that a sufficient condition for (i) and (ii) to hold is that, considering all K or less classes, the hessian of any interior point that satisfies the first-order conditions is definite negative (a formal proof would follow the same steps as those given in [10] for a similar property). We will refer to this sufficient condition as condition (S). Note that condition (S) is weaker than quasiconcavity. In fact, it is easy to construct instances under a *threshold policy* that do not have convex upper-level sets, hence are not quasiconcave, but do satisfy condition (S).

Let $\boldsymbol{\lambda}^* > 0$ be a feasible point of (P) that satisfies the first-order conditions (5). The hessian of the revenue function at $\boldsymbol{\lambda}^*$ can then be decomposed as:

$$H(W(\boldsymbol{\lambda}^*; \mathbf{C})) = \underbrace{\begin{bmatrix} \ddots & & 0 \\ & \frac{\bar{B}_i(\boldsymbol{\lambda}^*; \mathbf{C})}{\mu_i} \cdot R_i''(\lambda_i^*) & \\ 0 & & \ddots \end{bmatrix}}_{H_r} + \underbrace{\begin{bmatrix} \frac{\bar{B}_i(\boldsymbol{\lambda}^*; \mathbf{C})}{\mu_i} \cdot \frac{\bar{B}_j(\boldsymbol{\lambda}^*; \mathbf{C} - \mathbf{b}_i)}{\mu_j} \cdot \Delta_{ij} W(\boldsymbol{\lambda}^*; \mathbf{C}) \end{bmatrix}}_{H_p} \quad (12)$$

where $\Delta_{ij} W(\boldsymbol{\lambda}; \mathbf{C}) \equiv W(\boldsymbol{\lambda}; \mathbf{C}) - W(\boldsymbol{\lambda}; \mathbf{C} - \mathbf{b}_i) - W(\boldsymbol{\lambda}; \mathbf{C} - \mathbf{b}_j) + W(\boldsymbol{\lambda}; \mathbf{C} - \mathbf{b}_i - \mathbf{b}_j)$ is the double backward difference operator. Then, for condition (S) to hold, we must verify that (12) is

definite negative, considering all K classes or less.

PLACE FIGURE 4 HERE

When all R_i are strictly concave, clearly the diagonal matrix H_r is definite negative. The contribution of H_p will depend on the admission policy. In this sense, we can retrieve the results of the two previous subsections. Under *complete partitioning*, matrix H_p becomes a diagonal matrix with non-positive entries. Under *complete sharing*, $H_p = 0$ if $N = 1$, and if $N > 1$, then $H_p = (\overline{B}(\boldsymbol{\lambda}^*; N)\overline{B}(\boldsymbol{\lambda}^*; N-1)(\Delta W(\boldsymbol{\lambda}^*; N) - \Delta W(\boldsymbol{\lambda}^*; N-1))) \cdot [1/\mu_i\mu_j]$, which is definite negative, again due to the (strict) concavity of the nonblocking probability. However, the main result in this subsection is negative: *for general multi-class loss system operating under a general threshold policy, condition (S) does not hold*. In fact, consider a system with two classes and linear demands $\lambda_k(p_k) = \alpha_k - p_k\gamma_k$. The capacities of the common link and the outbound links are respectively $N = 2$, $N_1 = 2$, and $N_2 = 1$. Assuming, $\gamma_1 = 1/9$, $\alpha_1/\gamma_1 = 30$, $\gamma_2 = 1/7$, $\alpha_2/\gamma_2 = 0.02$, and $\mu_1 = \mu_2 = 1$, we can see in Figure 4 that the revenue function has an inflexion point and the level sets are not connected.

The reason for condition (S) not to hold in the previous example is that $N_2 = 1$ implies $\overline{B}_2(\boldsymbol{\lambda}^*; \mathbf{C} - \mathbf{b}_2) = 0$ so the lower right entry of H_p is zero, meaning that H_p can never be semi-definite negative. Then, we chose the problem parameters so that $H(W(\boldsymbol{\lambda}^*; \mathbf{C}))$ becomes indefinite.

It seems that having one of the outbound links with unit capacity is problematic. Note that the case $N = 1$ was always treated differently in the proofs of §4.3. The immediate question is: *does condition (S) hold if $N_k > 1 \forall k$?* Based on intensive numerical experiments (some reported in §6) we conjecture that the answer is affirmative when demands are linear.

5. Asymptotically Optimal Prices

In this section we consider a relaxation of the original static pricing problem (P). Suppose that the service provider is allowed to admit calls beyond the nominal link capacities, as long as the constraint is satisfied on average. Then he can charge static prices so that the average number of calls in the long run is below N in the common link and below N_k in each of the individual links. We call this relaxed problem P_{ub} . Since the average number of customers in an $M/M/\infty$ queue is equal to the traffic intensity, we have that problem P_{ub} can be formulated as the following nonlinear program:

$$\begin{aligned} (P_{ub}) \quad & \max \quad \sum_{k=1}^K R_k(\lambda_k)/\mu_k \\ & s.t. \quad \sum_{k=1}^K \lambda_k/\mu_k \leq N \\ & \quad \lambda_k/\mu_k \leq N_k \quad \forall k = 1 \dots K. \end{aligned} \tag{13}$$

By construction, the solution to the original optimization problem (P) satisfies the link capacity constraints on every sample path. Therefore, it is clear that the optimal value of the relaxed problem P_{ub} provides an upper bound for any static pricing policy. With the additional assumption that all the instantaneous revenue functions $R_k(\lambda_k)$ are concave, it can be shown that the arrival rates (or static prices) that solve the relaxed problem P_{ub} are asymptotically optimal for the original problem (P) (see [13, 14]). The asymptotic regime corresponds to scaling the capacities and the arrival rates by the same factor that tends to infinity, and can be seen as having *many small users*.

Let p_k^a be the static prices that solve P_{ub} . The performance of these prices outside the asymptotic regime is studied numerically in the next section. However, since the relaxed problem P_{ub} does not have call blocking, it is intuitively clear that the asymptotically optimal prices p_k^a will perform well in the original setting if they induce blocking probabilities that are close to zero. This fact is confirmed by the convergence rate proved in [14] which shows that system must be scaled until the blocking is zero for the asymptotic result to take place.

6. Numerical Results

In this section we provide some illustrative numerical examples with linear demands. We use the same notation introduced in the previous sections, with $\alpha_k = 10\gamma_k$, $\gamma_1 = 100$, $\gamma_2 = 20$, and $\mu_k = k$.

The first objective of this section is to assess the benefit (in terms of revenue) of solving (P) instead of using the asymptotically optimal prices. In other words, what is the loss in revenue from using p_k^a outside the limiting regime. Given the extra effort that might be required in solving (P) , there is a trade-off that can be relevant. We provide some insight in that direction.

We first consider a two-class *complete sharing* system. Since we assumed $\alpha_k = 10\gamma_k$, from the optimality conditions (5) it follows that the prices per minute are equal. This implies that the demand for class 1 is higher. In addition, class 1 calls stay longer in the system (on average), so class 1 customers seem to be more “valuable”. However, since $\beta_1 = \beta_2 = 10$, from §4.3 we know that both classes are “attractive” to the network service provider, and therefore we obtain an interior solution to problem (P) .

In Table 1, for increasing values of N , we compare the performance of the optimal static prices and those that are asymptotically optimal. The second and third columns come from $\rho_k = \alpha_k/N\mu_k$, and give a sense of the load offered by each class. Column four has the the average revenue obtained if the asymptotically optimal prices p_k^a are used as the static pricing policy (i.e. $W(\boldsymbol{\lambda}^a; \mathbf{C})$), and columns five and six have the optimal values of (P) and (P_{ub}) respectively. Similarly, columns seven and eight have the optimal prices in each case. Note that column seven confirms the monotonicity result of Proposition 4. Finally, as we anticipated at the end of the previous section, the last two columns show that when the blocking level goes down, so does the suboptimality gap of the asymptotic prices. In fact, these numerical results (together with others not reported here) suggest that the blocking probability must be less than 25% in order to have a suboptimality gap lower than 3%. In the case of the phone card company we worked with, such blocking level is realistic except for the peak hours during Sunday afternoon when a large number of expats are calling home.

PLACE TABLE 1 HERE

The next set of numerical experiments compares the performance of a two-class system over all possible *threshold policies* for a given capacity of the common link. That means, N is kept fixed (equal to five) and N_1 and N_2 range over $\{2, 3, 4, 5\}$. We excluded the cases $N_1 = 1$ or $N_2 = 1$ since, from the evidence of §4.4, the level sets of the revenue function might not be connected (and convergence to a global optimum is not guaranteed). For the remaining cases, we used third-party algorithms [18] to verify condition (S) (c.f. §4.4). The conclusion was that condition (S) holds in these cases, for *any* given set of parameters. Moreover, for these cases, if all α_k are equal, or all γ_k are equal, then condition (S) can be easily verified analytically through direct calculation of the

corresponding determinants (despite the extremely lengthy expressions the sign of the inequalities is trivially determined). These results, together with the same affirmative answer for many other instances, gives support to our belief that condition (S) might be satisfied in general under a *threshold policy*, linear demands, and $N_k > 1 \forall k$.

PLACE TABLE 2 HERE

Table 2 summarizes the numerical values obtained when the capacities of the outbound links vary. The demands and service requirements are the same as in Table 1. Again we observe that the revenue gap of the asymptotically optimal prices decreases as the blocking (for both classes) decreases. Another consistent and intuitive fact is that p_k^a is always greater than p_k^* , though note that the latter is not monotone with respect to increases in capacity.

7. Extensions and Conclusions

In this paper we presented a loss system model for the pricing problem faced by a network service provider. It was motivated by a company that sells phone cards.

We restricted our study to static pricing policies and focused on finding the optimal policy within that subclass. General first-order optimality conditions were derived, which have the same interpretation as many other revenue management problems: *the optimal prices must balance marginal revenue with opportunity costs*. An extensive analysis of second-order properties was done for two important subcases: *complete partitioning* and *complete sharing*. For these cases it was shown that the optimal static prices can be easily found using greedy search algorithms. In particular, determining the optimal static prices for the *complete sharing* case is as easy as maximizing a unimodal function.

The previous results do not extend to our general case (a *threshold policy*) but important intuition is given explaining the source of this negative outcome. The latter is strongly supported by numerical experiments. Finally, we addressed the question of when the asymptotic (*many relatively small users*) approximation should perform well. Our answer is that these prices are almost as good (in terms of revenue) as the optimal when blocking is not too high (i.e. less than 25% of the calls find the system busy).

We conclude this paper discussing which of our results extend to more general settings:

- **Demand substitution:** We assumed that the demand (arrival rate) of class k only depends on its own price. This dependence can be extended to the complete price vector to allow for demand substitution. Additional consistency assumptions, as those given in [4, 14], might be needed. Assuming that the demand function is invertible, the instantaneous rewards $R_k(\boldsymbol{\lambda}) = \lambda_k p_k(\boldsymbol{\lambda})$ are now a function of the arrival rate vector. The first-order conditions (for $\lambda_i^* > 0$) become:

$$\sum_{k=1}^K \frac{\partial R_k(\boldsymbol{\lambda}^*)}{\partial \lambda_i} \cdot \frac{\bar{B}_k(\boldsymbol{\lambda}^*)}{\mu_k} = \frac{\bar{B}_i(\boldsymbol{\lambda}^*)}{\mu_i} \cdot \left[W(\boldsymbol{\lambda}^*; \mathbf{C}) - W(\boldsymbol{\lambda}^*; \mathbf{C} - \mathbf{b}_i) \right] \quad (14)$$

Note that each term is weighted by the corresponding nonblocking probability and holding time, but the interpretation of “marginal revenue equal to opportunity cost” remains valid.

The results of §4.2 for the *complete partitioning* case no longer hold since the objective function is not separable. If $R_k(\boldsymbol{\lambda})$ is jointly concave for all k , then the proof of Proposition 3 for the *complete sharing* case is the same. Under this admission policy, the nonblocking probabilities in (14) cancel each other. Then, if the equations:

$$\sum_{k=1}^K \frac{\partial R_k(\boldsymbol{\lambda})}{\partial \lambda_i} \cdot \frac{\mu_i}{\mu_k} = \beta \quad i = 1, \dots, K,$$

can be solved uniquely for $\boldsymbol{\lambda}$ (as a function of β), then problem (P) can be reduced to one dimension as we did in §4.3.

An equation similar to (12) can be written for the demand substitution case in order to verify condition (S), but it is much more involved. Finally, the heuristic argument given in §5 on the suitability of the asymptotically optimal prices would still hold.

- **Heterogeneous resource requirements:** We examine what happens if class k has a bandwidth requirement of r_k , with not all of them equal. Clearly the first-order conditions (5) hold with the vector \mathbf{b}_i appropriately defined. The results for *complete partitioning* also follow, but those for *complete sharing* do not since the blocking probabilities are no longer equal.

The analysis of §4.4 for the *threshold policy* extends to the case of different resource requirements. In fact the equation for the hessian evaluated at an interior point that satisfies the the first-order conditions (cf. (12)) is valid. Hence, having $\mathbf{C} - \mathbf{b}_i > 0 \forall i$ is necessary for condition (S) to hold. Otherwise, $\bar{B}_i(\mathbf{C} - \mathbf{b}_i) = 0$ and counterexamples can be found. The discussion in §5 also extends to this case.

- **General product form loss network:** When the loss network has a general structure, as long as the stationary probability distribution has the product form (1), then the first-order conditions (5) will hold and the discussions in §4.4 and §5 will follow. However, calculating the normalization constant (1) for general loss networks is known to be an NP-complete problem and therefore might be a unsuccessful task. In those cases, solving approximated problems is the right path.
- **Games:** Extending the present model to a game theoretical framework when users act strategically is an open question and matter of future research.

References

- [1] Y. Almogly and O. Levin, A Class of Fractional Programming Problems, *Operations Research* 19(1) (1971) 57-67.
- [2] M. Avriel, W.E. Diewert, S. Schaible, and I. Zang, *Generalized Concavity*, (Plenum Publishing, New York, 1988).
- [3] E. Carrizosa, E. Conde, and M. Muñoz-Márquez, Admission Policies in Loss Queueing Models with Heterogeneous Arrivals, *Management Science*, 44(3) (1998) 311-320.

- [4] C.A. Courcoubetis and M.I. Reiman, Pricing in a Large Single Link Loss System, in: *Teletraffic Engineering in a Competitive World*, eds. P. Key and D. Smith, Elsevier, Amsterdam, 1999, pp. 737-746
- [5] C.A. Courcoubetis and R.R. Weber, *Pricing Communication Networks*, (John Wiley & Sons, West Sussex, England, 2003).
- [6] A. Ha., Optimal Pricing That Coordinates Queues with Customer-Chosen Service Requirements, *Management Science*, 47(7) (2003) 915-930.
- [7] A. Harel, Convexity Properties of the Erlang Loss Formula, *Operations Research*, 38(3) (1990) 499-505.
- [8] F.P. Kelly, Loss Networks, *The Annals of Applied Probability*, 1(3) (1991) 319-378.
- [9] S. Lanning, W.A. Massey, B. Rider, and Q. Wang, Optimal Pricing in Queueing Systems with Quality of Service Constraints, in: *Teletraffic Engineering in a Competitive World*, eds. P. Key and D. Smith, Elsevier, Amsterdam, 1999, pp. 747-757
- [10] V. Martínez de Albéniz and D. Simchi-Levi, Mean-Variance Trade-offs in Supply Contracts, Working Paper, MIT Operations Research Center, Cambridge, MA (2002).
- [11] H. Mendelson, Pricing Computer Services: Queueing Effects, *Communications of the ACM*, 28(3) (1985) 312-321.
- [12] H. Mendelson and S. Whang, Optimal Incentive-Compatible Priority Pricing for the M/M/1 Queue, *Operations Research*, 38(5) (1990) 870-883.
- [13] I.Ch. Paschalidis and J.N. Tsitsiklis, Congestion-Dependent Pricing of Network Services, *IEEE/ACM Transactions on Networking*, 8(2) (2000) 171-184.
- [14] I.Ch. Paschalidis and Y. Liu, Pricing in Multiservice Loss Networks: Static Pricing, Asymptotic Optimality, and Demand Substitution Effects, *IEEE/ACM Transactions on Networking*, 10(3) (2002) 425-438.
- [15] K.W. Ross, *Multiservice Loss Models for Broadband Telecommunication Networks*, (Springer-Verlag, New York, 1995).
- [16] K.W. Ross and D. Tsang, Algorithms to Determine Exact Blocking Probabilities for Multirate Tree Networks, *IEEE Transactions on Communications*, 38(8) (1990) 1266-1271.
- [17] K.W. Ross and D.D. Yao, Monotonicity Properties of the Stochastic Knapsack, *IEEE Transactions on Information Theory*, 36(5) (1990) 1173-1179.
- [18] A. Strzembonski, An Algorithm for Systems of Strong Polynomial Inequalities, *The Mathematica Journal*, 4(4) (1994) 74-77.
- [19] S. Ziya, H. Ayhan, and R.D. Foley, Optimal Pricing for a Service Facility, Working Paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, GA (2002).

N	ρ_1	ρ_2	J^a	J^*	J^{ub}	p_k^*	p_k^a	Blocking	$\frac{J^* - J^a}{J^*}$
2	500.00	50.00	11.98	18.81	19.96	9.70	9.98	0.94	36.33%
5	200.00	20.00	35.59	46.91	49.77	9.67	9.95	0.86	24.12%
10	100.00	10.00	77.83	93.40	99.09	9.63	9.91	0.76	16.67%
20	50.00	5.00	165.16	185.14	196.36	9.55	9.82	0.61	10.79%
30	33.33	3.33	253.16	275.22	291.82	9.47	9.73	0.50	8.01%
40	25.00	2.50	340.68	363.63	385.45	9.39	9.64	0.43	6.31%
50	20.00	2.00	427.26	450.39	477.27	9.31	9.55	0.37	5.14%
60	16.67	1.67	512.66	535.49	567.27	9.22	9.45	0.32	4.26%
70	14.29	1.43	596.75	618.93	655.45	9.14	9.36	0.28	3.58%
80	12.50	1.25	679.42	700.71	741.82	9.06	9.27	0.25	3.04%
90	11.11	1.11	760.61	780.84	826.36	8.98	9.18	0.23	2.59%

Table 1: Revenue gap: asymptotically optimal versus optimal (static prices).

N	N_1	N_2	J^a	J^*	J^{ub}	p_1^*	p_2^*	p_1^a	p_2^a	B_1	B_2	$\frac{J^* - J^a}{J^*}$
5	2	3	31.00	42.88	49.06	9.68	8.95	9.98	9.70	0.94	0.74	27.69%
5	2	4	32.95	43.15	49.06	9.75	8.98	9.98	9.70	0.93	0.72	23.64%
5	2	5	33.36	43.17	49.06	9.76	8.98	9.98	9.70	0.93	0.72	22.72%
5	3	2	31.32	44.34	49.51	9.68	9.02	9.97	9.80	0.91	0.82	29.37%
5	3	3	33.12	44.65	49.51	9.72	9.09	9.97	9.80	0.90	0.78	25.82%
5	3	4	33.58	44.69	49.51	9.73	9.10	9.97	9.80	0.90	0.77	24.85%
5	3	5	33.63	44.69	49.51	9.73	9.10	9.97	9.80	0.90	0.77	24.74%
5	4	2	33.78	45.89	49.74	9.69	9.25	9.96	9.90	0.88	0.84	26.40%
5	4	3	34.07	45.93	49.74	9.70	9.28	9.96	9.90	0.88	0.82	25.82%
5	4	4	34.10	45.94	49.74	9.70	9.28	9.96	9.90	0.88	0.81	25.76%
5	4	5	34.11	45.94	49.74	9.70	9.28	9.96	9.90	0.88	0.81	25.75%
5	5	2	35.57	46.90	49.77	9.67	9.68	9.95	9.95	0.86	0.87	24.17%
5	5	3	35.59	46.91	49.77	9.67	9.68	9.95	9.95	0.86	0.86	24.12%
5	5	4	35.59	46.91	49.77	9.67	9.67	9.95	9.95	0.86	0.86	24.12%
5	5	5	35.59	46.91	49.77	9.67	9.67	9.95	9.95	0.86	0.86	24.12%

Table 2: Static pricing and *threshold policies* for a two-class system.




Figure 1: A network service provider.

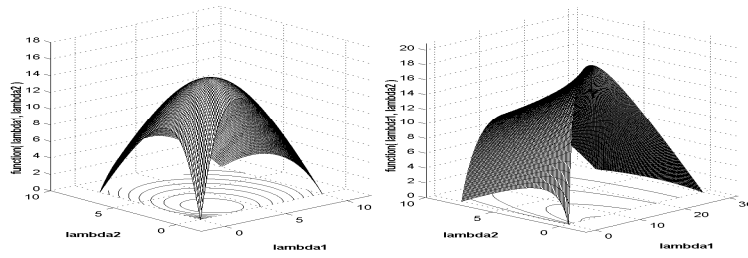


Figure 2: The revenue function for two classes with linear arrival rates.




Figure 3: The parametric functions $W(\lambda^P(\beta))$ and $\Delta W(\lambda^P(\beta))$.




Figure 4: Disconnected level sets for the *threshold policy* case.