

Augmenting Revenue Maximization Policies for Facilities where Customers Wait for Service

Avi Giloni

Syms School of Business, Yeshiva University, BH-428, 500 W 185th St., New York, NY 10033
agiloni@yu.edu

Yaşar Levent Koçağ

Syms School of Business, Yeshiva University, BH-428, 500 W 185th St., New York, NY 10033
kocaga@yu.edu

Phil Troy

Les Entreprises TROYWARE, 6590A Kildare, Cote St. Luc, Quebec, H4W 2Z4, Canada
Phil@PhilTroy.com

Abstract

In this paper, we study the dynamic pricing problem of a multi-server facility that processes requests from several customer classes on a first come first served basis. We assume an arrival belongs to one of a finite number of customer classes and that each class is distinguished by known and arbitrary service valuations and arbitrary but non-decreasing waiting cost functions. We model the facility as an $M/M/S/I$ queue and use the theory of Markov Decision Processes to identify dynamic pricing strategies that maximize the revenues obtained from customers who are assumed to be acting individually to maximize their utility. The key to our approach is recognizing, in context of our model, that the maximum revenue obtainable from a service facility is bounded by the maximum collective benefits less waiting costs that customers could receive in such a facility which is achieved through state dependent social (welfare) optimization. For our model, we show that this upper limit can be achieved with state dependent revenue maximization when it is possible to exactly charge customers the benefit they receive less the waiting costs they incur. This in turn can be accomplished when it is possible to identify the group of customers to which each customer belongs, where customers in each group have the same benefit and waiting cost function. We further demonstrate that some or most of this additional revenue can also be achieved even when customers in each group only have similar benefits and waiting cost functions. Finally, we illustrate that the revenue maximizing pricing policy might actually increase arrival rates as system occupancy increases, by admitting customer classes that were not admitted in previous states.

Keywords: Revenue Management, State Dependent Pricing, Dynamic Pricing, Markov Decision Process, Queueing, Delay Sensitive Customers.

1. Introduction

Since Naor (1969), there has been a considerable amount of research on managing congestion in facilities modeled as queues. While Naor focussed on determining the socially or organizationally optimal policy for such facilities, i.e. the policy that maximizes the collective benefits less waiting costs that users or internal customers receive, this research has since expanded to include the maximization of revenues generated by external customers.

In this paper, we augment state dependent revenue maximization approaches for facilities that process customer requests on a first come first served basis. The idea behind our approach is that when it is possible to identify the group to which a customer belongs, it may also be possible to charge group specific tolls and thus increase revenues. The key to our approach is recognizing, in context of our model, that the maximum revenue obtainable from a service facility is bounded by the maximum collective benefits less waiting costs that internal customers could receive in such a facility. For our model, we show that this upper limit can be achieved with state dependent revenue maximization when it is possible to exactly charge customers the benefit they receive less the waiting costs they incur. This in turn can be accomplished when it is possible to identify the group of customers to which each customer belongs, where customers in each group have the same benefit and waiting cost function. We further demonstrate that some or most of this additional revenue can also be achieved even when customers in each group only have similar benefits and waiting cost functions.

To accomplish this, we develop a generic state dependent control model that we subsequently use to analyze and compare an organizationally optimal policy and several revenue maximization policies. Using that model we show that applying our approach, when the firm can identify groups of customers having exactly the same benefits and waiting cost functions, yields revenues that are at least as large as those obtainable when the firm can only identify groups of customers with similar benefits and waiting cost functions, which in turn yields revenues at least as large as when the firm cannot differentiate customers. We further show that our approach also results in a control policy that generally makes it possible for the firm to take advantage of its ability to differentiate customers in the presence of competitors that do not.

Though limited to facilities that serve customers on a first come first served basis, the results of our analysis suggest that our approach might also be used to augment policies involving state independent control mechanisms, more general service time distributions and other scheduling mechanisms that provide higher priority to customers with higher waiting costs, as well as those

using different computational approaches such as diffusion or fluid models.

1.1 Organization

Before starting our analysis, in Section 1.2, we briefly review existing research on revenue management, state dependent social optimization, state dependent revenue maximization, and the integration of revenue maximization with production optimization. Then, in order to make it possible to evaluate the efficacy of our approach relative to existing organizational optimization and revenue maximization, we start our analysis in Section 2 by specifying a generic service facility model, and its associated control problem, in context of an M/M/S/I queuing model. Because that control problem is a Markov Decision Process, we specify its policy iteration value determination equations. From those equations, we identify opportunity costs inherent to the problem and determine general circumstances under which these opportunity costs increase in the number of jobs in the facility. Having this generic model model as a platform, in the remainder of the paper we apply it to the different control problems that we study. Thus in section 3, we study the state dependent social optimization control policy (P1). In section 4, we study the state dependent revenue maximization control policy (P2). In section 5, we study the state and super-group dependent revenue maximization control policy (P3).

In particular, in section 3 we formulate and characterize the state dependent social optimization control policy. In so doing, we extend Naor's (1969), Miller's (1969) and Whang's (1986) analysis to handle multiple servers and multiple groups of customers, where each group has its own arrival rate, job benefit, and arbitrary but non-decreasing waiting cost function. The most important contribution of this section is that we show for our model that the rate at which state dependent social optimization generates customer benefits less waiting costs is an upper bound for the rate at which any state dependent policy based on our model can generate revenues. To gain insight into how this policy works, we show that it ensures that only jobs whose benefit less expected waiting costs exceeds the opportunity costs caused by their admission are admitted to the facility and that these opportunity costs increase in the number of jobs in the facility. As a result, the groups of customers whose jobs are admitted for processing when there is a specific number of jobs in the facility is a subset of those admitted when there are less jobs in the facility. This in turn implies that the rate at which jobs are accepted for processing is non-increasing and typically decreasing in the number of jobs in the facility. Should it be necessary or desired, we demonstrate that tolls, set to these state dependent opportunity costs, can be used to implement this policy.

We proceed in section 4 by adapting the analysis of section 2 to state dependent revenue

maximization without requiring, as is done elsewhere, that the rate at which revenue is generated via arrivals to be continuously differentiable or have a downward slope. In so doing, we extend existing analysis of state dependent revenue maximization to handle multiple servers and multiple groups of customers, where each group has its own arrival rate, job benefit, and arbitrary waiting cost function. The most important result of this section is that we show that Chen and Frank's (2001) result that state dependent revenue maximization cannot fully capture customer benefits less waiting costs when customers are heterogeneous still holds, even when a non-discounted model is used. As we did for state dependent social optimization, we show that the state dependent opportunity costs associated with state dependent revenue maximization increase in the number of jobs in the facility. In contrast, we show that state dependent revenue maximization does not admit the same jobs as does social optimization because revenue maximization only admits jobs to the facility if they maximize the rate at which state dependent revenues, less their corresponding opportunity costs, are generated. As an extension to Whang (1986), we prove that when customers all have the same waiting cost function, the optimal arrival rate decreases in the number of jobs in the facility. However, because this is not a simple threshold policy, we also show that, unlike social optimization, when groups of customers have different waiting cost functions there can be situations in which jobs from groups excluded in one state are processed in a higher state. When that happens it is possible for the optimal job arrival rate to increase as the number of jobs in the facility increases.

To address the limitation that state dependent revenue maximization cannot fully capture customer benefits less waiting costs when customers are heterogeneous, in section 5 we propose and characterize a state and super-group dependent revenue maximization policy that adds customer differentiation to state dependent revenue maximization. Because this policy takes advantage of a facility's ability to identify the super-group, i.e. the group of groups, to which a customer belongs, it makes it possible to capture more customer net benefits than does the state dependent revenue maximization policy. We show that when each super-group contains exactly one group that this policy yields exactly the same admission control as does state dependent social optimization, but fully captures customer benefits less waiting costs as revenues. Finally, in section 6 we verify, through examples, the efficacy of each policy, and in section 7 we summarize our results and suggest future work.

1.2 Literature Review

The pioneering work of Naor(1969) started a stream of literature, commonly known as congestion pricing, which studies pricing decisions in service systems modeled as queues where customers are rational decision makers who maximize the utility they derive from the service by explicitly considering the customer disutility due to queueing delays. In his work, Naor considers an $M/M/1$ queue and shows that it can be socially suboptimal to allow customers to decide whether to submit requests for processing as doing so can cause an externality in the form of additional waiting for subsequent customers. Naor's model is extended by Yechiali (1971) to a $GI/M/1$ queue, by Knudsen (1972) to an $M/M/c$ queue, and by Lippman and Stidham (1977) to a general birth and death process with bounded, non-decreasing and concave departure rates. Mendelson (1985) further extends this analysis by incorporating capacity decisions for facilities serving internal customers who only see average queue lengths and have linear delay costs. In follow-up papers, Dewan and Mendelson (1990) allow for non-linear delay cost functions while Mendelson and Whang (1990) incorporate multiple customers classes and incentive-compatible priority pricing mechanisms. Notice that in all the aforementioned papers the pricing decision is static while we address dynamic pricing decisions that adapt to the level of congestion in the system.

There is also a vast literature on dynamic control of arrivals to queues that focuses on pricing. We review only those papers which are closely related to ours and refer the interested reader to the survey papers by Bitran and Caldentey (2003), Elmaghraby (2002) and McGill and van Ryzin (1999). Miller (1969) studies a Markovian multi-server loss system where the rewards for each customer class is fixed and admission is controlled dynamically to maximize average expected rewards over an infinite horizon. Feinberg and Yang (2011) extend the admission control problem studied in Miller (1969) to a multi-server system with finite waiting room and fixed customer rewards that also depend on the level of congestion in the system and establish the optimality of threshold policies with respect to long run average profit as well as several other optimality criteria. Paschalidis and Tsitsiklis (2000) also study a Markovian multi-server loss system and identify the dynamic pricing policy that maximizes long term average expected revenue. Low (1974a) and Low (1974b) consider the dynamic pricing of an $M/M/s$ multiserver queue with finite and infinite waiting rooms, respectively. He restricts attention to a finite action space where each action corresponds to a particular price which in turn results a particular arrival rate. He further assumes that the system incurs bounded holding costs and shows that there exists a stationary policy where prices increase as the system gets more congested. Maoui *et al* (2007) extend this setting to one where

holding costs are unbounded and customer behavior is explicitly captured through willingness-to-pay distributions. All the aforementioned papers (except for Feinberg and Yang(2011) who study admission control) either disregard the customer disutility due to waiting or assume that it is the facility that incurs these costs. Our work on the other hand assumes several customer classes that are also differentiated by their sensitivity to delay and therefore yields some interesting insights unique to our setting.

Next, we discuss some closely related dynamic pricing papers in greater detail. In his unpublished paper, Whang (1986) formulates the state dependent revenue maximization problem for a service facility modeled as an $M/M/1$ queue in which all external customers have the same linear waiting cost function and a concave revenue rate function. He shows that the optimal policy decreases job submission rates as the number of jobs in the facility increases. While our result's echo Whang (1986) when all customers have the same waiting cost, we also show that when the waiting cost functions are different, it is possible that the job submission rate can actually increase. Furthermore, we allow for multiple servers, an arbitrary but finite waiting space, arbitrary service valuations as well as arbitrary but non-decreasing waiting cost functions.

Ata and Shenorson (2005) formulate, solve, and characterize the socially optimal policy for simultaneously controlling job acceptance and processing rates of a service facility that can be modeled as an $M/M/1$ queue. The most important contribution of this work is that it integrates social optimization with production management. Assuming that all internal users have the same linear waiting cost function and that capacity costs are non-decreasing, convex, and zero when capacity is zero, they show that the optimal control policy results in decreasing job acceptance rates and increasing processing rates as the number of jobs waiting for processing increases, even though the optimal tolls are not necessarily monotonic in the number of jobs in the facility. We note that our work is different in that our focus is on revenue maximization and that we allow for a discrete number of customer classes which are also allowed to have arbitrary yet non-decreasing waiting cost functions.

Chen and Frank (2001) formulate and analyze the revenue maximization problem for an $M/M/1$ facility where external customers have the same linear waiting cost function and receive one of two discrete job benefits. They show that when using state dependent revenue maximization with external customers having homogeneous benefits that the firm will serve the same set of customers as the social planner but that when external customers have heterogeneous benefits that firms cannot capture all available surplus as revenues. Notice that in contrast to their work which maximizes long run discounted revenues for facilities modeled as a single server queue, our work

focuses on long run average revenue for facilities modeled as single or multiple server queues. Due to our model formulation, we are directly able to observe that the maximum revenue of the facility is bounded by the social optimal and that the facility is able to achieve this solution under perfect discrimination which is in line with the observation in Chen and Frank (2001). Finally, because we allow for arbitrary waiting cost functions for different classes, we are able to see interesting results unique to our setting such as the observation that the arrival rate to the facility might increase with congestion.

2. The Model

We start our analysis by describing a generic service facility model that we subsequently use throughout the paper; because it is generic, this facility model has no explicit objective function. In context of this model, we then formulate a generic control problem that we subsequently adapt to both state dependent social optimization and state dependent revenue maximization for comparative purposes. To gain insights into this model, we identify opportunity costs that vary with the number of jobs in the facility and that are subsequently used to characterize state dependent social optimization and state dependent revenue maximization. We also identify conditions, used in those characterizations, that show when these opportunity costs increase in the number of jobs in the facility.

2.1 Model Specification

In our model, we assume that K groups of customers bring jobs to a service facility for processing. Each group belongs to a unique super-group having similar job processing benefits and waiting cost functions; super-groups are identified as a_1, a_2, \dots, a_C , where $C \leq K$. Group k customers arrive at the facility at an average rate of λ_k per unit time with exponentially distributed interarrival times. If they submit their job and it is accepted for processing, group k customers receive a benefit b_k after job processing is completed. While waiting for processing to be completed, group k customers incur a waiting cost of $w_{i,k}$ whenever there are i jobs ahead of them. We assume that $w_{i,k}$ is non-decreasing in i and that $\lim_{i \rightarrow \infty} w_{i,k} = \infty$ for all k . We define $\beta_{i,k} = b_k - w_{i,k}$ to be the net benefit that group k customers receive if their jobs are admitted to the facility when the facility is already in state i . Customers only submit jobs for processing if their net benefit is greater than or equal to any toll they are charged and balk otherwise.

To facilitate analysis, we model our service facility as an $M/M/S/I$ queue having S servers and

a buffer that can hold $I - S$ jobs¹. The facility processes all accepted jobs on a first come first served basis at the rate of μ_i jobs per unit time where μ is the average number of jobs processed by each server per unit time and

$$\mu_i = \begin{cases} i\mu & \text{if } 1 \leq i < S \\ S\mu & \text{otherwise.} \end{cases} \quad (1)$$

2.2 State Dependent Optimization Via Policy Iteration

In context of this model, we formulate a generic state dependent control problem that we subsequently adapt to both state dependent social optimization and state dependent revenue maximization. We observe that the problem is a Markov decision process, that it can be formulated as a non-discounted continuous time policy iteration problem, and that the optimal solution can be characterized directly from that formulation. To do so, we specify the value determination equations for such problems (Howard 1960),

$$\gamma = r_{i,i} + \sum_{\iota \neq i} a_{i,\iota} r_{i,\iota} + \sum_{\iota} a_{i,\iota} v_{\iota}. \quad (2)$$

In context of our problem and under a particular policy

- γ is the expected reward generated per unit time,
- $r_{i,i}$ is the rate at which reward is generated while in state i ,
- $r_{i,\iota}$, for $\iota \neq i$, is the reward received when the facility changes state from i to ι ,
- $a_{i,\iota}$, for $\iota \neq i$, is the rate at which the facility changes state from i to ι ,
- $a_{i,i} = -\sum_{\iota} a_{i,\iota}$,
- v_i is the relative benefit of being in state i .

We further define

- $\rho_{i,k}$ as the reward received when a group k job is admitted when the facility is in state i and
- $\xi_{i,k}$ as the fraction of group k jobs accepted when the facility is in state i . $\xi_{i,k}$ must either be the control parameter, or be controllable by the control parameter. It must also always be possible, either directly or indirectly, to set $\xi_{i,k}$ to 0 so as to ensure that there are no admissions of group k customers when the facility is in state i .

¹Note that because $\lim_{i \rightarrow \infty} w_{i,k} = \infty$ for all k , there exists a finite $N := \min(i : \beta_{ik} \leq 0 \text{ for all } k=1,2,\dots,K)$ after which every customer will choose to balk. Thus, the restriction to a finite queue capacity I is without loss of generality.

We observe that there is no reward for staying in state i and that transitions are only between state i and $i - 1$ and between state i and $i + 1$. We also observe that the rate of transition between state i and $i + 1$ is $\sum_k \xi_{i,k} \lambda_k$ and that the rate of transition between state i and $i - 1$ is μ_i . Thus, the value determination equations are

$$\gamma = \sum_k \xi_{i,k} \lambda_k \rho_{i,k} + \sum_k \xi_{i,k} \lambda_k v_{i+1} + \mu_i v_{i-1} - \sum_k \xi_{i,k} \lambda_k v_i - \mu_i v_i. \quad (3)$$

To simplify (3), we define $\nabla v_i = v_i - v_{i+1}$ and obtain

$$\gamma = \sum_k \xi_{i,k} \lambda_k (\rho_{i,k} - \nabla v_i) + \mu_i \nabla v_{i-1}. \quad (4)$$

Given these value determination equations, the policy iteration algorithm for the completely ergodic continuous time decision processes (c.f. Howard 1960) is to arbitrarily pick an initial policy, and to iteratively, until γ converges for all states, solve the linear value determination equations for ∇v_i and find the policy that maximizes γ for each value of i using the current values of ∇v_i . Finally, we note that our optimality equations have a solution and that the policy iteration algorithm will indeed converge to the optimal solution. This is because it can easily be shown that our problem satisfies the first recurrency condition in Federgruen and Tijms (1978). We omit the formal proof for brevity.

2.3 ∇v_i are Opportunity Costs

Critical to an understanding of the model is the observation that ∇v_i is the opportunity cost of admitting a job to the facility when it is in state i . To see this, consider (4) applied to state I , defined to be the highest state the facility attains under a particular policy, for which there are no job admissions and for which $\mu_I = S\mu$ so that

$$\gamma = S\mu \nabla v_{I-1} \text{ which implies that } \nabla v_{I-1} = \frac{\gamma}{S\mu}. \quad (5)$$

We recall that γ is the expected reward generated per unit time and that $S\mu$ is the expected number of jobs processed per unit time when the system is in state I making transitions to state $I - 1$. Thus ∇v_{I-1} is the expected reward per job of making transitions from state I to state $I - 1$, or equivalently, the expected cost, and hence the opportunity cost, of making transitions from state $I - 1$ to I by admitting jobs when the facility is in state $I - 1$. Repeating this unit analysis for lower states we see that ∇v_i is the opportunity cost of admitting a job to the facility when the facility is in state i .

2.4 ∇v_i are Increasing in i

To facilitate the characterization of the state dependent social optimization and state dependent revenue maximization control policies, we show that the state dependent opportunity costs increase in i under conditions that we subsequently show hold for both state dependent social optimization and state dependent revenue maximization. We start by defining

$$m_i = \max \sum_k \xi_{i,k} \lambda_k (\rho_{i,k} - \nabla v_i)$$

where the maximization is with respect to the control variables of the current policy, given known values of ∇v_i for the optimal solution. We recall that the policy improvement step requires the maximization of the value determination equations, again given known values of ∇v_i . Thus at optimality, equation 4 can be restated as

$$\gamma = m_i + \mu_i \nabla v_{i-1}. \quad (6)$$

We subtract adjacent instances of this equation and rearrange terms to obtain the result that at optimality

$$\mu_{i+1} \nabla v_i - \mu_i \nabla v_{i-1} = m_i - m_{i+1}. \quad (7)$$

In what follows, we provide three conditions which will allow us to prove Theorem 1 and subsequently establish the monotonicity of opportunity costs associated with an optimal state dependent policy under social optimization and revenue maximization in Sections 3 and 4, respectively.

Condition 1. *If $m_i \leq m_{i+1}$, then $\nabla v_{i+1} \leq \nabla v_i$ for all $0 \leq i \leq I$.*

Condition 2. *If $m_i > m_{i+1}$, then $\nabla v_{i+1} > \nabla v_i$ for $i < S - 1$.*

Condition 3. *If $\nabla v_{i+1} > \nabla v_i$, then $m_i > m_{i+1}$ for $i \geq S - 1$.*

We begin by establishing the following lemma which implies that ∇v_0 cannot be non-negative while Condition 1 holds.

Lemma 1. *Suppose Condition 1 holds and $\nabla v_0 \leq 0$. Then the facility generates a non-positive rate of reward.*

Proof. Let $\nabla v_0 \leq 0$. Consider (7) when $i = 0$ for which the left hand side of (7) is non-positive since $\mu_0 = 0$ as there are no jobs to process. Thus, the right hand side is non-positive as well, which implies that $m_i \leq m_{i+1}$, which in turn implies that $\nabla v_1 \leq \nabla v_0$ because of condition 1. We next

consider (7) where $i = 1$ for which the left hand side is again non-positive since $\mu_{i+1} \geq \mu_i$. Thus, the right hand side is again non-positive which implies that $\nabla v_2 \leq \nabla v_1$. Repeating this analysis for $i > 1$, we see that $\nabla v_{I-1} \leq \nabla v_{I-2} \leq \dots \leq \nabla v_1 \leq \nabla v_0 \leq 0$. This along with (5) implies that the service facility generates a non-positive rate of reward. \square

Using this lemma, we now prove that the state dependent opportunity costs increase in state under Conditions 1-3.

Theorem 1. *When the conditions 1-3 hold, the opportunity costs associated with an optimal state dependent policy associated with (4) are increasing in state for those states in which the facility admits customers, in the presence of arbitrary but discrete benefits, multiple servers, and customers having heterogeneous and non-linear waiting cost functions.*

Proof. By Lemma 1, ∇v_0 can only be positive. By considering (7), condition 2 of of Lemma 1 and analysis similar to that used in that lemma's proof, it is easy to show that

$$\nabla v_{S-1} > \nabla v_{S-2} > \dots > \nabla v_1 > \nabla v_0. \quad (8)$$

We observe for $i \geq S$ that (7) can be restated as

$$S\mu\nabla v_i - S\mu\nabla v_{i-1} = m_i - m_{i+1} \quad (9)$$

and that for $i = I - 1$, (9) becomes

$$S\mu\nabla v_{I-1} - S\mu\nabla v_{I-2} = m_{I-1} \quad (10)$$

since there are no admissions to the facility when the facility is in its highest state. As m_{I-1} by definition must maximize the rate of reward less opportunity costs the facility generates when in state $I - 1$, and as the control policy allows the facility to disallow admissions, the facility will only admit customers in state $I - 1$ when they generate a positive rate of reward. This implies that the right hand side of (10) is positive, which in turn implies that $\nabla v_{I-1} > \nabla v_{I-2}$ since $S\mu(\nabla v_{I-1} - \nabla v_{I-2}) > 0$. We consider (9) for $S < i = I - 2$ and note that since $\nabla v_{I-1} > \nabla v_{I-2}$ and by condition 3 that, $m_i > m_{i+1}$. Thus in order for (9) to hold for the case where $i = I - 2$, $\nabla v_{I-2} > \nabla v_{I-3}$. It follows that for $S \leq i < I - 2$ that $\nabla v_i > \nabla v_{i-1}$. This combined with (8) completes the proof. \square

3. State Dependent Social Optimization: P1

In this section, we formulate and characterize the state dependent social optimization control policy and show that it provides an upper bound on the performance of the revenue maximization policies developed in subsequent sections.

3.1 Problem Formulation

To determine the socially optimal policy for a service facility, we observe that the social reward generated for making a transition between state i and $i + 1$ equals $\beta_{i,k}$, the net benefit customers receive when their jobs are accepted for processing. Thus, the value determination equations for the state dependent socially optimal policy are

$$\gamma^s = \sum_k \xi_{i,k}^s \lambda_k (\beta_{i,k} - \nabla v_i^s) + \mu_i \nabla v_{i-1}^s, \quad (11)$$

where the superscript s indicates that the objective under consideration is socially optimality.

3.2 The Optimal Policy

From Section 2.2, we know that to implement the policy improvement step for (11), we maximize γ^s with respect to the control variables $\xi_{i,k}^s$ for each value determination equation, using the current values of ∇v_i^s . Since the $\mu_i \nabla v_{i-1}^s$ term is irrelevant to this maximization, we only need to maximize $\sum_k \xi_{i,k}^s \lambda_k (\beta_{i,k} - \nabla v_i^s)$ with respect to the control variables $\xi_{i,k}^s$. Since there is a control parameter for each group and since the terms for each group are separable, we can maximize $\xi_{i,k}^s \lambda_k (\beta_{i,k} - \nabla v_i^s)$ with respect to $\xi_{i,k}^s$ separately for each group. In doing so, we see that when

- $\beta_{i,k} > \nabla v_i^s$, this expression is maximized by setting $\xi_{i,k}^s$ to 1, i.e. by admitting group k jobs
- $\beta_{i,k} < \nabla v_i^s$, this expression is maximized by setting $\xi_{i,k}^s$ to 0, i.e. by not admitting group k jobs
- $\beta_{i,k} = \nabla v_i^s$, this expression is maximized by setting $\xi_{i,k}^s$ to any value in $[0, 1]$ and it doesn't matter whether group k jobs are admitted.

In other words, the optimal policy maximizes the rate at which facilities generate net benefits by only admitting jobs whose net benefits are greater than or equal to any opportunity costs they create.

For organizations that use internal charge-back mechanisms, having already assumed that customers only submit jobs whose net benefit is greater than or equal to any toll they are charged, we observe that this policy can be implemented by charging state dependent tolls set to ∇v_i^s .

3.3 Characterization of The Optimal Policy

To provide a better understanding of this policy, we prove that the state dependent opportunity costs that underlie it increase in the number of jobs in the facility. We show that this property

implies that the groups of customers whose jobs are processed in state $i+1$ are a subset of the groups of customers whose jobs are processed in state i . This in turn implies that the total acceptance rate of jobs for processing by the facility in state i is non-increasing and typically decreasing in i . We also prove that the rate at which revenues can be generated under any state dependent revenue maximization policy is bounded by the rate at which state dependent social optimization generates net benefits.

We start by proving that the state dependent opportunity costs that underlie state dependent social optimization increase in the number of jobs in the facility. To simplify the proof, we define k_i^s to be the set of groups whose jobs are admitted to the facility when it is in state i under state dependent social optimization. By definition of k_i^s and the analysis of Section 3.2, $\ell \in k_i^s$ if $\xi_{i,\ell}^s = 1$. We also define $\Lambda_i = \sum_{k \in k_i^s} \lambda_k$ to be the total rate at which jobs are accepted when the facility is in state i .

Lemma 2. *The state dependent socially optimal policy guarantees that*

- a) *For $1 \leq i \leq I$, if $\nabla v_i^s \leq \nabla v_{i+1}^s$, then $k_{i+1}^s \subseteq k_i^s$ which implies that $\Lambda_{i+1} \leq \Lambda_i$.*
- b) *For $i < S - 1$, if $\nabla v_i^s > \nabla v_{i+1}^s$, then $k_i^s \subseteq k_{i+1}^s$ which implies that $\Lambda_i \leq \Lambda_{i+1}$.*

Proof. To prove a), we recall that under state dependent social optimization that if $k \in k_i^s$ then $\beta_{i,k} \geq \nabla v_i^s$. Since $\beta_{i+1,k} \leq \beta_{i,k}$ because expected waiting costs cannot decrease in i , if $\nabla v_i^s \leq \nabla v_{i+1}^s$, it follows that if $\beta_{i+1,k} \geq \nabla v_{i+1}^s$ then $\beta_{i,k} \geq \nabla v_i^s$. Thus k_{i+1}^s cannot include any groups that are not in k_i^s and thus Λ_{i+1} must be less than or equal to Λ_i . To prove b), we first observe that for $i < S - 1$ that $\beta_{i+1,k} = \beta_{i,k}$ since arriving customers do not have to wait for processing to begin in these states. Then if $\nabla v_i^s > \nabla v_{i+1}^s$, $\beta_{i,k} > \nabla v_i^s$ implies that $\beta_{i+1,k} > \nabla v_{i+1}^s$ and k_i^s cannot include any groups that are not in k_{i+1}^s and thus Λ_i must be less than or equal to Λ_{i+1} . \square

Theorem 2. *The opportunity costs that arise under state dependent social optimization increase in state for states in which the facility admits customers, in the presence of arbitrary but discrete benefits, multiple servers and customers having heterogeneous and non-linear waiting cost functions.*

Proof. By Theorem 1, to prove this theorem we only need to show that the conditions of Lemma 1 hold. To show that condition 1 holds, we assume that $m_i \leq m_{i+1}$. This implies that there exists some k such that $\xi_{i,k} \lambda_k (\beta_{i,k} - \nabla v_i) \leq \xi_{i+1,k} \lambda_k (\beta_{i+1,k} - \nabla v_{i+1})$. This along with the fact that $\beta_{i,k} \geq \beta_{i+1,k}$ imply that $\nabla v_{i+1} \leq \nabla v_i$ and thus condition 1 holds. To show that condition 2 holds, we assume that $i < S - 1$ and that $m_i > m_{i+1}$. This implies that there exists some k

such that $\xi_{i,k}\lambda_k(\beta_{i,k} - \nabla v_i) > \xi_{i+1,k}\lambda_k(\beta_{i+1,k} - \nabla v_{i+1})$. This along with the fact that $\beta_{i,k} = \beta_{i+1,k}$ imply that $\nabla v_{i+1} > \nabla v_i$ and thus condition 2 holds. To show that condition 3 holds, we assume that $\nabla v_{i+1}^s > \nabla v_i^s$. Without restriction of generality, we assume that group k customers are admitted to the facility when it is in state i . This along with the fact that $\beta_{i,k} \geq \beta_{i+1,k}$ imply that $\beta_{i,k} - \nabla v_i > \beta_{i+1,k} - \nabla v_{i+1}$. Since group k customers are admitted, it follows that $m_i > m_{i+1}$ and thus condition 3 holds. Thus, the theorem is proven. \square

Since ∇v_i and thus the optimal tolls are increasing in i , we know from Lemma 2 that the groups of customers whose jobs are accepted for processing in state $i + 1$ are a subset of those accepted for processing in state i . We also know from Lemma 2 that this implies that the total rate at which jobs are accepted is non-increasing and typically decreasing in i . We observe that the reason these phenomena occur is because of the simple nature of state dependent social optimization's control mechanism which admits all jobs whose net benefit is greater than a threshold value. The result of these phenomena is that the optimal control policy tends to keep the facility busy while minimizing queue lengths and waiting costs.

Besides being important on its own account, this analysis also suggests that the rate at which revenues can be generated for our model under any state dependent revenue maximization policy is bounded by the rate at which state dependent social optimization generates customer net benefits. To see that this is indeed the case, we recall that the maximum toll customers are willing to pay for processing equals their net benefit and that the rate at which net benefits can be generated is maximized by social optimization. Thus the upper bound of the rate that customers will pay for processing equals the rate at which social optimization generates net benefits.

4. State Dependent Revenue Maximization: P2

Having just shown that the upper bound of the rate that customers will pay for processing equals the rate at which social optimization generates net benefits, in this section we develop and analyze a state dependent revenue maximization policy in which a single toll is selected for each state so as to maximize the rate at which revenues are generated. Through this analysis, we demonstrate how state dependent revenue maximization differs from state dependent social optimization and thus why and when it generates less revenue than the upper bound provided by state dependent social optimization. Notice also that our model formulation allows us to give a clearer comparison between state dependent social optimization and state dependent revenue maximization and thus gives deeper insights into the differences between the two than that has been done already in the literature.

4.1 Problem Formulation

To specify the value determination equations for state dependent revenue maximization, we observe that the reward generated under such a policy comes from the tolls, θ_i , specific to each state, charged to customers submitting jobs when the system is in state i . Thus, the value determination equation for state dependent revenue maximization for each state is

$$\gamma^r = \sum_k \xi_{i,k}^r \lambda_k (\theta_i - \nabla v_i^r) + \mu_i \nabla v_{i-1}^r, \quad (12)$$

where the superscript r indicates that the objective is revenue maximization.

In contrast to their role under state dependent social optimization, $\xi_{i,k}^r$ are not control variables, but are instead controlled by the state dependent tolls. As discussed in Section 2, customers submit jobs when the net benefit they receive for processing of those jobs is greater than or equal to any tolls they are charged for processing. Thus, $\xi_{i,k} = I(\beta_{i,k} \geq \theta_i)$, where $I(\cdot)$ is the indicator function.

4.2 The Optimal Policy

To determine the optimal tolls, θ_i , we apply the policy improvement step to (12) by maximizing γ^r with respect to θ_i . Since the $\mu_i \nabla v_{i-1}^r$ terms are irrelevant to this maximization, we need only maximize

$$\sum_k \xi_{i,k}^r \lambda_k (\theta_i - \nabla v_i^r) \quad (13)$$

with respect to θ_i , which we can rewrite as

$$\sum_k \lambda_k I(\beta_{i,k} \geq \theta_i) (\theta_i - \nabla v_i^r). \quad (14)$$

From (14), we note that policy improvement will always ensure that $\theta_i \geq \nabla v_i^r$. This is because $\theta_i < \nabla v_i^r$ will lead to negative values of (14) which policy improvement can preclude by setting θ_i to a larger value.

In order to apply the policy improvement step, since we do not assume differentiability or downward sloping revenue rate functions, we search for the value of θ_i that maximizes (14) for each state. We observe that this search need only be over the values of $\beta_{i,k}$ that are greater than ∇v_i^r . To see that this is true, consider a candidate toll, $\theta > \nabla v_i^r$, not equal to one of the $\beta_{i,k}$, which results in some arrival rate. We observe that if we were to increase this toll to the smallest $\beta_{i,k}$ greater than this toll that the arrival rate would not change and that the total revenues would increase. Thus, θ_i must equal one of the $\beta_{i,k}$ or without restriction of generality ∇v_i^r if $\nabla v_i^r > \beta_{i,k}$ for all k .

To gain some insights into the optimal policy, let us assume that the optimal toll for state i , $\theta_i^* = \beta_{i,\ell} > \nabla v_i^r$. In such a case, (14) becomes

$$\sum_k \lambda_k I(\beta_{i,k} \geq \beta_{i,\ell}) (\beta_{i,\ell} - \nabla v_i^r). \quad (15)$$

Thus, for state i , there exist three category of customers: customers who do not submit their job to be processed because their net benefit after toll would be negative (\mathbf{G}_i), customers who submit jobs to be processed but have no surplus, i.e., their net benefit is equal to the toll that is charged (\mathbf{E}_i), and customers who submit jobs to be processed and have surplus (\mathbf{L}_i). We denote $\mathbf{K} = \{1, \dots, k\}$ as the set of indices representing all customer groups. Thus, we have $\mathbf{G}_i = \{j \in \mathbf{K} | \beta_{i,j} < \beta_{i,\ell}\}$, $\mathbf{E}_i = \{j \in \mathbf{K} | \beta_{i,j} = \beta_{i,\ell}\}$, and $\mathbf{L}_i = \{j \in \mathbf{K} | \beta_{i,j} > \beta_{i,\ell}\}$. The surplus that customer group $j \in \mathbf{L}_i$ receives is equal to $\beta_{i,j} - \beta_{i,\ell}$. Hence, the firm can only capture all benefits less waiting costs as revenues under revenue maximization if all customers who submit jobs have the same benefits less waiting costs.

4.3 Characterizing the Opportunity Costs and Arrival Rates

To provide a better understanding of this policy, we prove that the state dependent opportunity costs that underlie it increase in the number of jobs in the facility. We then observe that because this policy does not result in the acceptance of all jobs whose net benefit is greater than these state dependent opportunity costs, that it is possible for a group not admitted when the facility is in state i to be admitted when the facility is in state $i + 1$. Using these results, we show that when all customer groups have the same waiting cost function, the state dependent revenue maximization policy leads to total arrival rates in state i that are decreasing in i . However, when customer groups do not all have the same waiting cost function this relationship may not hold, thus further differentiating the control aspects of this policy from state dependent social optimization.

We start by proving that the state dependent opportunity costs that underlie state dependent revenue maximization increase in the number of jobs in the facility.

Theorem 3. *The revenue maximizing state dependent opportunity costs, ∇v_i^r , are increasing in state, for those states that the facility is admitting customers, in the presence of arbitrary but discrete benefits, multiple servers and customers having heterogeneous and possibly non-linear waiting cost functions.*

Proof. To prove the theorem, we need only show that the three conditions from Lemma 1 hold. To show that condition 1 holds, assume that $m_i \leq m_{i+1}$. In this proof, we assume that θ_i maximizes

$\sum_k I(\beta_{i,k} > \theta_i) \lambda_k(\theta_i - \nabla v_i^r)$. Since θ_i is assumed to maximize $\sum_k I(\beta_{i,k} > \theta_i) \lambda_k(\theta_i - \nabla v_i^r)$ and $\beta_{i,k} \geq \beta_{i+1,k}$, it follows that

$$\begin{aligned} m_i &= \sum_k I(\beta_{i,k} > \theta_i) \lambda_k(\theta_i - \nabla v_i^r) \geq \sum_k I(\beta_{i,k} > \theta_{i+1}) \lambda_k(\theta_{i+1} - \nabla v_i^r) \\ &\geq \sum_k I(\beta_{i+1,k} > \theta_{i+1}) \lambda_k(\theta_{i+1} - \nabla v_i^r). \end{aligned} \quad (16)$$

Now, assume that $\nabla v_i^r < \nabla v_{i+1}^r$. This implies that

$$\sum_k I(\beta_{i+1,k} > \theta_{i+1}) \lambda_k(\theta_{i+1} - \nabla v_i^r) > \sum_k I(\beta_{i+1,k} > \theta_{i+1}) \lambda_k(\theta_{i+1} - \nabla v_{i+1}^r) = m_{i+1} \quad (17)$$

which is a contradiction and thus condition 1 follows.

To show that condition 2 holds, we assume that $i < S-1$ and $m_i > m_{i+1}$. Since θ_{i+1} maximizes m_{i+1} and $\beta_{i,k} = \beta_{i+1,k}$, it follows that

$$m_{i+1} \geq \sum_k I(\beta_{i+1,k} > \theta_i) \lambda_k(\theta_i - \nabla v_{i+1}^r) = \sum_k I(\beta_{i,k} > \theta_i) \lambda_k(\theta_i - \nabla v_{i+1}^r). \quad (18)$$

Now assume that $\nabla v_{i+1}^r \leq \nabla v_i^r$. This along with (18) implies that

$$m_{i+1} \geq \sum_k I(\beta_{i,k} > \theta_i) \lambda_k(\theta_i - \nabla v_{i+1}^r) \geq m_i \quad (19)$$

which is a contradiction, and thus condition 2 holds.

To show that condition 3 holds, we assume that $\nabla v_{i+1}^r > \nabla v_i^r$. This along with the fact that θ_i maximizes m_i and $\beta_{i,k} \geq \beta_{i+1,k}$ implies that

$$\begin{aligned} m_i &\geq \sum_k I(\beta_{i,k} > \theta_{i+1}) \lambda_k(\theta_{i+1} - \nabla v_i^r) \geq \sum_k I(\beta_{i+1,k} > \theta_{i+1}) \lambda_k(\theta_{i+1} - \nabla v_i^r) \\ &> \sum_k I(\beta_{i+1,k} > \theta_{i+1}) \lambda_k(\theta_{i+1} - \nabla v_{i+1}^r) = m_{i+1}. \end{aligned}$$

Thus, condition 3 holds and the theorem is proven. \square

Unlike state dependent social optimization, this result, when combined with the observation that the optimal policy is not a threshold policy, suggests the possibility that a group not admitted in state i may be admitted in state $i+1$. This can occur because state dependent revenue maximization selects tolls that implicitly determine those groups that will submit jobs to the facility in order to maximize revenues less opportunity costs for each state. Thus when waiting costs substantially differ between groups it is possible that groups that submit jobs in state $i+1$ may not do so in state i . Thus in contrast to state dependent social optimization, it is possible for group k to submit jobs in state $i+1$ even though the toll was too high for them to do so in state i .

This observation leads to the following theorem.

Theorem 4. *When all customer groups have the same waiting cost function, state dependent revenue maximization leads to a total arrival rate Λ_i that is non-increasing in the state i . When customer groups do not all have the same waiting cost function, state dependent revenue maximization does not necessarily result in a total arrival rate Λ_i that is non-increasing in the state i and the total arrival rate Λ_i might actually increase for some state i .*

Proof. To prove the first part of the theorem, we recall that ∇v_i^r increases in state i and that $\beta_{i,k}$ does not increase in state i . Let $\theta_i^*(\Lambda)$ be the toll in state i needed to obtain a job submission rate of Λ , let Λ_i^* be the optimal job submission rate for state i and let Λ be any other job submission rate that is greater than Λ_i^* . Because Λ_i^* is the optimal job submission rate we know that

$$\Lambda(\theta_i(\Lambda) - \nabla v_i^r) - \Lambda_i^*(\theta_i(\Lambda_i^*) - \nabla v_i^r) \leq 0. \quad (20)$$

Now consider state $i + 1$. Let Δw_i be the increase in expected waiting costs and $\Delta \nabla v_i$ be the increase in ∇v_i^r , both of which must be non-negative. We observe that

$$\Lambda(\theta_{i+1}(\Lambda) - \nabla v_{i+1}^r) = \Lambda(\theta_i(\Lambda) - (\Delta w_i + \nabla v_i^r + \Delta \nabla v_i)) \quad (21)$$

Thus by (20),

$$\Lambda(\theta_{i+1}(\Lambda) - \nabla v_{i+1}^r) - \Lambda_i^*(\theta_{i+1}(\Lambda_i^*) - \nabla v_{i+1}^r) \quad (22)$$

$$= \Lambda(\theta_i(\Lambda) - (\Delta w_i + \nabla v_i^r + \Delta \nabla v_i)) - \Lambda_i^*(\theta_i(\Lambda_i^*) - (\Delta w_i + \nabla v_i^r + \Delta \nabla v_i)) \quad (23)$$

$$\leq (\Lambda - \Lambda_i^*)(-\Delta w_i - \Delta \nabla v_i) < 0 \quad (24)$$

where the last inequality follows because Δw_i and $\Delta \nabla v_i$ are both non-negative and since we are only considering Λ greater than Λ_i^* . As a result, a job submission rate $\Lambda > \Lambda_i^*$ cannot be maximizing $\Lambda_i(\theta_i - \nabla v_i)$ which in turn implies that the optimal arrival rate Λ_i^* has to be non-increasing in the state i .

Next, we prove the second part of this theorem by considering an example in which there is one server and two groups of customers with different waiting costs. Let $\lambda_1 < \lambda_2$, and $b_1 > b_2$. We assume that the waiting costs of the two groups are such that $\beta_{i,1} > \beta_{i,2}$ for $i = 0, \dots, S$, and $\beta_{i,1} \leq 0 < \beta_{N,2}$ for $i = S + 1, \dots, I$. We further assume that

$$\lambda_1 \beta_{i,1} \geq (\lambda_1 + \lambda_2) \beta_{i,2} \quad \text{for } i = 0, \dots, S. \quad (25)$$

and without loss of generality that the policy iteration algorithm starts with an initial policy in which $\theta_i = \beta_{i,1} \geq \nabla v_i$ for $i = 0, \dots, S$ and $\theta_i > \beta_{i,2}$ for all i . As a result, group 1 customers submit

jobs to the facility when the facility is in states 0 through S and no customers submit jobs in higher states. We observe that $\nabla v_S = \gamma/s\mu$ and that it can be easily calculated for this situation. We apply the policy improvement step, maximizing $\Lambda_i(\theta_i - \nabla v_i)$, and from (25) see that it is always suboptimal to set $\theta_i = \beta_{i,2}$ for $i = 0, \dots, S$. However, if $\beta_{S,2} \geq \nabla v_{S-1}$, then the policy improvement step sets $\theta_S = \beta_{S,2}$. Should that happen, the policy improvement step would never set $\theta_S > \beta_{S,2}$ as otherwise the algorithm would cycle and not converge. Therefore, we see that

$$\Lambda_i = \begin{cases} \lambda_1 & 0 \leq i \leq S-1 \\ \lambda_2 & i = S. \end{cases} \quad (26)$$

Since $\lambda_2 > \lambda_1$, the arrival rate is indeed increasing and the theorem is proven. \square

5. State and Super-group Dependent Revenue Maximization: P3

In an attempt to determine whether it is possible to improve upon state dependent revenue maximization policy, in this section we introduce a policy that takes advantage of any ability the facility has to identify the super-group, i.e. group of groups, to which a customer belongs. In this state and super-group dependent revenue maximization policy, there is a different toll for each combination of state and super-group.

Of particular importance, in this section we show that this new policy generates at least as much, and generally more, revenue than does state dependent revenue maximization. We further show that when each super-group contains exactly one group that the admissions of this policy are equivalent to that of state dependent social optimization, but that in contrast to state dependent social optimization all net benefits are captured as tolls.

5.1 Problem Formulation

In contrast to state dependent revenue maximization in which there is only one toll for each state, in this policy there is a toll, θ_{i,a_c} , for each combination of state and super-group; when there is exactly one group in each super-group there is a toll for each state and group combination. Thus, the value determination equations for state and super-group dependent revenue maximization are

$$\gamma^a = \sum_{c=1}^C \sum_{k \in a_c} \xi_{i,k}^a \lambda_k (\theta_{i,a_c} - \nabla v_i^a) + \mu_i \nabla v_{i-1}^a, \quad (27)$$

where the superscript a indicates that the objective is state and super-group dependent revenue maximization and where $\xi_{i,k}^a$ equals 1 if $\beta_{i,k}$ is greater than or equal to θ_{i,a_c} and 0 otherwise when group k is in super-group a_c .

5.2 The Optimal Policy

To maximize revenues, we apply the policy improvement step to (27) by maximizing γ^a with respect to θ_{i,a_c} . Since the $\mu_i \nabla v_{i-1}^a$ terms are irrelevant to this maximization, we need only maximize

$$\sum_{c=1}^C \sum_{k \in a_c} \xi_{i,k}^a \lambda_k (\theta_{i,a_c} - \nabla v_i^a) \quad (28)$$

with respect to θ_{i,a_c} and because each super-group has its own control variable we can do this separately for each super-group.

Following the analysis of Section 4.2, we observe that for states in which the net benefit is less than ∇v_i^a for all groups in a super-group, that the toll for that super-group will equal ∇v_i^a and no customers from that super-group will submit jobs to the facility. On the other hand, if there are groups in a super-group whose net benefit is greater than ∇v_i^a then the optimal toll for that super-group is selected to maximize

$$\sum_{k \in a_c} \xi_{i,k}^a \lambda_k (\theta_{i,a_c} - \nabla v_i^a). \quad (29)$$

Further following the analysis of Section 4.2, the optimal toll will always equal one of the $\beta_{i,k}$ of the groups in that super-group. In the event that the facility can identify the specific group to which each customer belongs so that the number of super-groups equals the number of groups, then the optimal toll will always equal $\beta_{i,k}$ when it is greater than or equal to ∇v_i^a .

5.3 Characterization of The Optimal Policy

To provide a better understanding of this policy, we first observe that the state dependent opportunity costs that underlie it increase in the number of jobs in the facility. We next show that this policy generates at least as much revenue, and generally more, than does state dependent revenue maximization. We also consider the situation in which the number of groups equals the number of super-groups and show that the resulting admission control is identical to that of state dependent social optimization and that the revenue generated equals the net benefit that would be generated under state dependent social optimization.

To prove that the state dependent opportunity costs that arise under this policy increase in state, one can adapt the proof of Theorem 3 to allow for super-groups which we skip for brevity. To show that this policy always generates at least as much revenue, and generally more, than does state dependent revenue maximization, we apply policy iteration starting with the state dependent revenue maximizing tolls as the initial policy. We observe that by setting the state and super-group

dependent tolls equal to the state dependent revenue maximizing tolls, the value determination equations for this policy are equivalent to those of state dependent revenue maximization. We apply the policy improvement step with the state and super-group dependent tolls. If it is impossible to improve the policy for any of the states, then policy iteration is done and the resulting policy is identical to and generates the same revenues as does state dependent revenue maximization. Should it be possible to improve the policy for even one state, then the resulting policy will be better (c.f. Howard (1960)). For it to be impossible to improve the policy for at least one state, it would require that there be no super-groups unadmitted under state dependent revenue maximization having any groups that have net-benefits greater than the existing opportunity costs. It would also require that there be no super-groups whose optimal toll under this policy is higher than the corresponding toll under state dependent revenue maximization. Since the likelihood of both of these conditions holding is relatively low when there is more than one super-group, we see that this policy always generates at least as much revenue as does state dependent revenue maximization, and generally more.

Finally, we consider the situation in which the number of groups equals the number of super-groups. When this occurs, the value determination equations are equivalent to

$$\gamma^a = \sum_{c=1}^C \sum_{k \in a_c} \xi_{i,k}^a \lambda_k (\beta_{i,k} - \nabla v_i^a) + \mu_i \nabla v_{i-1}^a, \quad (30)$$

since the optimal toll for each super-group will either equal $\beta_{i,k}$ or ∇v_{i-1}^a , and since when the latter situation occurs $\xi_{i,k}^a$ will equal 0. Thus the value determination equations when the number of groups equals the number of super-groups are equivalent to those of state dependent social optimization, and the upper bound on revenues generated discussed in Section 3.3 is achieved.

These results point to the fact that when it is possible to identify customers' super-groups, this policy makes it possible to remove much or all of the inefficiencies associated with state dependent revenue maximization. Often, the ability to identify the super-group and/or group to which a customer belongs is available due to market research activities. In the event that this information is known by the firm, the super-group policy makes it possible for managers to extract additional revenues, since it only requires an ability to identify the super-group to which a customer belongs, rather than the specific group. If this information is available only external to the firm and at a cost, the firm can calculate the value of this information by comparing the rate of revenue of the two models in order to determine the value of the information.

6. Examples

In this section, we verify through two examples the efficacy of each policy. For each example, we first compute the maximum rate of net benefit obtainable via state dependent social optimization. We then compute the revenue generated by the other policies and compare them against this maximum. To show the effect of having just one group per super-group, we consider an example having two groups per super-group and compare it with a similar example in which there is only one group per super-group.

6.1 Example 1: 2 Super-Groups Each Having 2 Groups

We consider an example in which the service facility has 2 servers each of which processes jobs at the rate of 5 jobs per unit time, and allows up to 25 jobs to be in the facility at one time². We first assume that there are two groups in each of two super-groups. In the first super-group, group 1 customers arrive at the rate of 2.5 jobs per unit time, receive a benefit for job processing of 600, and incur a waiting cost of $w_{i1} = 10 [(i - 2)^+]^2$ when in they find the system in state i with $(i - 2)^+$ customers already waiting in queue. Group 2 customers receive a benefit for job processing of 400 but otherwise have the same characteristics. In the second super-group, group 3 customers arrive at the rate of 1.5 jobs per unit time, have a gross job benefit of 150, and incur a waiting cost of $w_{i2} = (i - 2)^+$ when they find the system in state i with $(i - 2)^+$ customers already waiting in queue. Group 4 customers receive a benefit of 100 for job processing but otherwise have the same characteristics.

As predicted in the previous section, in Table 1 we first observe that rate at which net benefit is generated under state dependent social optimization (P1), 1621, is greater than the rate at which revenue is generated under all the other policies since each of the super-groups has more than one group³. We next observe that the rate at which revenue is generated under state dependent revenue maximization (P2), 797, is approximately fifty percent less than the rate at which net benefit is generated under state dependent social optimization. We also observe that the rate at which revenue is generated under the state and super-group dependent optimization policy (P3), 1314, is almost but not quite as large as the rate at which net benefit is generated under state dependent social optimization. This is because there is more than one group per super-group.

²Following our initial comment regarding the buffer size, one might argue that it could be beneficial to consider the admission of group 3 and group 4 customers even after state 25. However, the infinite horizon expected average revenue, γ , is insensitive to the buffer choice as long as the size of the buffer is large enough and the choice of 25 is without loss of generality.

³We note that this value along with the associated infinite horizon average expected revenues (γ) values of other (P2) and (P3) are reported as Δv_{25} in Table 1 and Table 2.

Thus it can be very advantageous to use state and super-group dependent revenue maximization when it is possible to identify the super-group of each customer.

i	P1				P2				P3			
	Δv_i	θ_i	θ_i	Λ_i	Δv_i	θ_i	θ_i	Λ_i	Δv_i	θ_i	θ_i	Λ_i
0	16	16	16	8.00	1	400	400	2.00	11	400	100	8.00
1	26	26	26	8.00	5	400	400	2.00	18	400	100	8.00
2	49	49	49	8.00	24	400	400	2.00	33	400	100	8.00
3	74	74	74	8.00	30	99	99	8.00	48	390	99	8.00
4	97	97	97	8.00	35	98	98	8.00	67	360	148	5.00
5	122	122	122	5.00	41	97	97	8.00	83	310	147	5.00
6	143	143	143	5.00	48	96	96	8.00	99	440	146	4.00
7	154	154	154	1.00	55	95	95	8.00	115	350	145	4.00
8	162	162	162	1.00	59	94	94	7.00	127	240	144	4.00
9	162	162	162	0.00	64	93	93	7.00	127	10000	143	3.00
10	162	162	162	0.00	65	92	92	6.00	127	10000	142	3.00
11	162	162	162	0.00	67	91	91	6.00	128	10000	141	3.00
12	162	162	162	0.00	68	90	90	6.00	128	10000	140	3.00
13	162	162	162	0.00	69	89	89	6.00	129	10000	139	3.00
14	162	162	162	0.00	71	88	88	6.00	129	10000	138	3.00
15	162	162	162	0.00	72	87	87	6.00	130	10000	137	3.00
16	162	162	162	0.00	74	86	86	6.00	130	10000	136	3.00
17	162	162	162	0.00	75	85	85	6.00	130	10000	135	3.00
18	162	162	162	0.00	76	84	84	6.00	131	10000	134	3.00
19	162	162	162	0.00	76	83	83	6.00	131	10000	133	3.00
20	162	162	162	0.00	76	82	82	6.00	131	10000	132	3.00
21	162	162	162	0.00	75	81	81	6.00	131	10000	10000	0.00
22	162	162	162	0.00	71	80	80	6.00	131	10000	10000	0.00
23	162	162	162	0.00	65	79	79	6.00	131	10000	10000	0.00
24	162	162	162	0.00	80	128	128	3.00	131	10000	10000	0.00
25	1621	1621	1621	0.00	797	10000	10000	0.00	1314	10000	10000	0.00

Table 1: Characterization of Optimal Dynamic Pricing Policies for Example 1

6.2 Example 2: 2 Super-Groups Each Having 1 Group

To see the effect of having just a single group in each super-group, we modify the previous example so that the groups in each super-group are merged and their job benefits are averaged. Thus, group 1 customers arrive at the rate of 2 jobs per unit time, have a gross job benefit of 500, and incur a waiting cost of $w_{i1} = 10 [(i - 2)^+]^2$ when they find the system in state i . Group 2 customers arrive at the rate of 6 jobs per unit time, have a gross job benefit of 125, and incur a waiting cost of $w_{i2} = (i - 2)^+$ when they find the system in state i . The facility has 2 servers, processes jobs at the rate of 5 jobs per unit time, and allows up to 25 jobs to be in the facility at one time.

As in the previous example, in Table 2 we first observe that rate at which net benefit is generated under state dependent social optimization (P1), 1613, is greater than or equal to the rate at which revenue is generated under all the other policies. We next observe that that the rate at which revenue is generated under state dependent revenue maximization, 997 (P2), is approximately forty percent less than the rate at which net benefit is generated under state dependent social optimization. In contrast to the above example, we see that the rate at which revenue is generated under the state and super-group dependent optimization policy exactly equals the rate at which net benefit is generated under state dependent social optimization (P3). This is because there is exactly one group per super-group. Finally, we see that in both Table 1 and Table 2 that state dependent revenue maximization (P2) results in a decrease in the posted price and an increase in the total arrival rate in state 3.

i	P1				P2				P3			
	Δv_i	θ_i	θ_i	Λ_i	Δv_i	θ_i	θ_i	Λ_i	Δv_i	θ_i	θ_i	Λ_i
0	17	17	17	8.00	1	500	500	2.00	17	500	125	8.00
1	28	28	28	8.00	5	500	500	2.00	28	500	125	8.00
2	52	52	52	8.00	27	500	500	2.00	52	500	125	8.00
3	79	79	79	8.00	33	124	124	8.00	79	490	124	8.00
4	104	104	104	8.00	40	123	123	8.00	104	460	123	8.00
5	122	122	122	2.00	47	122	122	8.00	122	410	10000	2.00
6	144	144	144	2.00	55	121	121	8.00	144	340	10000	2.00
7	161	161	161	2.00	64	120	120	8.00	161	250	10000	2.00
8	161	161	161	0.00	75	119	119	8.00	161	10000	10000	0.00
9	161	161	161	0.00	76	118	118	6.00	161	10000	10000	0.00
10	161	161	161	0.00	78	117	117	6.00	161	10000	10000	0.00
11	161	161	161	0.00	79	116	116	6.00	161	10000	10000	0.00
12	161	161	161	0.00	81	115	115	6.00	161	10000	10000	0.00
13	161	161	161	0.00	82	114	114	6.00	161	10000	10000	0.00
14	161	161	161	0.00	84	113	113	6.00	161	10000	10000	0.00
15	161	161	161	0.00	85	112	112	6.00	161	10000	10000	0.00
16	161	161	161	0.00	87	111	111	6.00	161	10000	10000	0.00
17	161	161	161	0.00	88	110	110	6.00	161	10000	10000	0.00
18	161	161	161	0.00	90	109	109	6.00	161	10000	10000	0.00
19	161	161	161	0.00	91	108	108	6.00	161	10000	10000	0.00
20	161	161	161	0.00	93	107	107	6.00	161	10000	10000	0.00
21	161	161	161	0.00	94	106	106	6.00	161	10000	10000	0.00
22	161	161	161	0.00	96	105	105	6.00	161	10000	10000	0.00
23	161	161	161	0.00	98	104	104	6.00	161	10000	10000	0.00
24	161	161	161	0.00	100	103	103	6.00	161	10000	10000	0.00
25	1613	1613	1613	0.00	997	10000	10000	0.00	1613	10000	10000	0.00

Table 2: Characterization of Optimal Dynamic Pricing Policies for Example 2

7. Conclusions and Future Work

In this paper, we demonstrated that it is possible to augment state dependent revenue maximization to take advantage of capabilities that facilities have of differentiating between groups of customers. Specifically, we showed that when it is possible to differentiate customers, it is often possible to charge different tolls for each type of customer and thus increase revenues.

To do so, we extended and integrated existing work on state dependent social optimization, state dependent revenue maximization, and product differentiation. In particular, we extended existing work on state dependent social optimization and revenue maximization to simultaneously handle multiple servers and multiple customer groups, where each customer group has its own job benefit and non-decreasing but otherwise arbitrary waiting cost function. We then introduced and characterized a state and super-group dependent revenue maximization policy that takes advantage of facilities' ability to differentiate customers. Because this policy takes advantage of facilities' ability to identify the super-group, i.e. the group of groups, to which a customer belongs, it makes it possible to capture more customer net benefits than does the state dependent revenue maximization policy. We also showed that when it is possible to identify each customer's group, this policy exactly yields the same admission control as does state dependent social optimization while fully capturing customers' benefits less waiting costs as revenues.

Our work suggests several open research problems. These include (i) The adaptation of this approach to allow for more general service time distributions and for mechanisms that schedule customers based on their waiting cost functions such as those developed by Afeche and Mendelson (2004) and Mendelson and Whang (1990); (ii) The adaptation of this approach to the diffusion and fluid approaches developed by Maglaras (2006); and (iii) The adaptation of our work to that of Ata and Shenorson (2005) so as to maximize revenues using state and super group revenue maximization, while also adjusting capacity in a state dependent manner.

References

- Afeche, P., and H. Mendelson. 2004. Pricing and Priority Auctions in Queueing Systems with a Generalized Delay Cost Structure. *Management Science*. **50** (7) 869-882.
- Ata, B., and S. Shneorson. 2005. Dynamic Control of an M/M/1 Service System with Adjustable Arrival and Service Rates. *Management Science*. **52**(11) 1778-1791.
- Bitran, G., and R. Caldentey. 2003. An Overview of Pricing Models for Revenue Management. *MSOM* **11**(3) 203-229.

- Chen, H., and M. Z. Frank. 2001. State Dependent Pricing with a Queue. *IIE Transactions*. **33** 847-860.
- Elmaghraby, W., and P. Keskinocak. 2003. Dynamic Pricing: Research overview, current practices and future directions. *Management Science*. **49**(10) 1287-1309.
- Federgruen A., and Tijms, H. C. 1978. The optimality equation in average cost denumerable state semi-Markov decision problems, recurrency conditions and algorithms. *Journal of Applied Probability*. **15** 356–373.
- Feinberg E. A., and Yang, F. 2011. Optimality of Trunk Reservation for an M/M/k/N Queue with Several Customer Types and Holding Costs. *Probability in the Engineering and Informational Sciences*, published online 21 July 2011 DOI:10.1017/S0269964811000179 .
- Howard, R. 1960. *Dynamic Programming and Markov Processes* The MIT Press.
- Knudsen, N. C. 1972. Individual and Social Optimization in a multi-server Queue with a General Cost-Benefit Structure. *Econometrica*. **40**(3) 515-528.
- Lippman, S.A., and S. Stidham 1977. Individual versus Social Optimization in Exponential Congestion Facilities. *Operations Research*. **25**(2) 233-247.
- Low, D. W. 1974. Optimal dynamic pricing policies for an $M/M/s$ queue. *Operations Research*. **22**(3) 545-561.
- Low, D. W. 1977. Optimal pricing for an unbounded queue. *IBM Journal of Research and Development*. **18**(4) 290-302.
- Maglaras, C. 2006. Revenue management for a multiclass single-server queue via a fluid model analysis. *Operations Research*. **54**(5) 914-932.
- McGill, J. I., and G. J. van Ryzin. 1999. Revenue Management: Research overview and prospects. *Transportation Science*. **33**(2) 233-256.
- Mendelson, H., and S. Whang. 1990. Optimal Incentive-Compatible Pricing for the M/M/1 Queue. *Management Science*. **38**(5) 870-883.
- Mendelson, H. 1985. Pricing Computer Services: Queueing Effects. *Communications of the ACM*. **28**(3) 312-321.
- Miller, B.L. 1969. A Queueing Reward System With Several Customer Classes. *Management Science*. **16**(3) 234-245.
- Naor, P. 1969. The regulation of queue size by levying tolls. *Econometrica*. **37**(1) 15-24.

- Paschalidis, I.C. and Tsitsiklis, J.N. 2000. Congestion-dependent pricing of network services. *IEEE/ACM Transactions on Networking* . **8**(2) 171184.
- Whang, S. 1986. The Value of System State Information in Information Systems, working paper.
- Yechiali, U. 1971. On Optimal Balking Rules and Toll Charges in the GI/M/1 Queueing Process, *Operations Research*. **19**2 349-370.