Optimal Information Disclosure Policies in Strategic Queueing Games

Eran Simhon[a,*], Yezekael Hayel[b,c], David Starobinski[a], Quanyan Zhu[c]

[a]*Division of Systems Engineering, Boston University, 8 Saint Marys St., Boston, MA 02215, USA*
[b]*LIA/CERI, University of Avignon, 339 chemin des Meinajaries cedex 9, 84911 Avignon, France*
[c]*Department of Electrical and Computer Engineering, Tandon School of Engineering, New York University, 5 MetroTech Center, Brooklyn, NY 11201,USA*

Abstract

Information about queue length is an important parameter for customers who face the decision whether to join a queue or not. In this paper, we study how optimal information disclosure policies can be used by a service provider of an $M/M/1$ queue to increase its revenue. Our main contribution is showing that the intuitive policy of informing customers about the current queue length when it is short and hiding this information when it is long is never optimal.

Keywords: Queueing games, Partial information system, Revenue optimization.

1. Introduction

The literature on the strategic behavior of customers in queues is traditionally divided into the classical observable and unobservable queues. In the former case, introduced by Naor [1], customers are informed about the current queue length before deciding whether to join or balk. In the latter case, introduced by Edelson and Hilderbrand [2], customers make their decisions based on statistical information (e.g., the queue parameters). In both cases, customers behave strategically and join the queue only if their expected waiting cost is smaller than the reward obtained upon being served.

The goal of this work is to find out whether there are situations where the service provider can increase its revenue by combining those two frameworks. In particular, we are interested in studying policies where the provider shares information with some customers and hides it from others, depending on the actual queue length. We assume that the service provider has a fixed income from each customer that joins the queue. Thus, in order to maximize its revenue, the provider should maximize the effective arrival rate, which is the rate of customers that join the queue, or equivalently, minimize the idle period of its system.

We compare between the outcomes of the following three policies: (*i*) always inform customers about the queue length; (*ii*) never inform customers about the queue length; (*iii*) inform customers based on a threshold policy in which queue length information is provided when the queue length is below the specified threshold and is hidden otherwise. Although the third policy seems intuitive, we formally prove that, in any setting, either sharing information with all customers or hiding information from all customers always yields greater expected revenue.

The rest of the paper is structured as follows. In Section 2, we discuss related works. In Section 3, we formally define the model. In Section 4, we find the equilibrium structure of a game with threshold information policy. In Section 5, we compare between the outcomes of the three policies. In Section 6, we prove our main result.

2. Related Works

Information disclosure plays a key factor in customers' behavior at call centers. For a recent review of call center theory with many references in the role of information on the balking behavior of customers see [3]. The authors of [4] give examples of queueing systems providing on-line queue length information, including emergency rooms in hospitals, voting locations, security gates at international airports and amusements parks. Most of these systems assume that providing queue-length helps customers to optimize their expected payoffs.

Both [5] and [6] compare between the outcomes of observable and unobservable $M/M/1$ queues. It is

---

*Corresponding author.
    Email address:* `simhon@bu.edu` (Eran Simhon)

shown that each model can outperform the other depending on the settings. In [7], the authors consider a population with two types of customers: informed and uninformed. The proportion of informed customers is an exogenous parameter of the model. They analytically characterize the equilibrium and find that the effective incoming rate is unimodal with respect to the fraction of informed customers. A similar result is presented in [4], where the model assumes that each customer follows a two-step decision process: (i) to buy or not to buy the information about the queue; (ii) to join or not join the queue. The characterization of an equilibrium strategy is based on an analysis of topological properties of the expected utility sets. The authors prove the existence and uniqueness of a symmetric Nash equilibrium.

In our work, we adopt a control-theoretic perspective. In particular, we are interested in finding the best information disclosure policy that should be used by a service provider in order to maximize its revenue. The fact that the disclosed information is state-dependent complicates the analysis of the equilibrium. In [8], the authors consider informed and uninformed customers, but the latter are non-strategic and follow a determined threshold policy. In [9], the authors consider customers' decisions in observable and unobservable queueing games with endogenously determined arrivals and batch services. The equilibrium strategies are evaluated numerically.

3. System Model

We consider a standard $M/M/1/$ FCFS queue system. The arrival rate is a Poisson process with mean λ. The service rate is exponential with mean $1/\mu$. We denote $\rho = \lambda/\mu$ to be the *maximum* load of the queue (recall that customers do not always join the queue). The cost of each time unit spent at the queue (waiting and being served) is C. Without loss of generality we set $C = 1$. All customers have the same reward R from service, where $R > 1/\mu$ (otherwise, no customer ever joins the queue).

The standard set-up in all the problems related to strategic behavior in queuing systems [10] is that a new customer decides to join or to balk depending on her expected sojourn time in the queue and the reward obtained for service completion. The behavior of a new customer is highly dependent on her knowledge/information about the current queue length (the queue length at the instant of her arrival). Naturally, a customer will join the queue if and only if the expected sojourn time (which depends on the information about

the queue length provided to her) is smaller than the reward.

We consider a state-dependent information disclosure policy, such that the control parameter u is a mapping from the queue length i into the interval $[0, 1]$. Then, the control $u(i)$ for all $i = 0, 1, \ldots$ is the probability that the provider gives the information to an arrival customer when there are i customers in the queue. We only consider stationary information disclosure policies, i.e. control policies that do not depend on time. If $u(i) = 1$, for all $i \geq 0$, then it is exactly the (fully) observable model presented by Naor [1] and if $u(i) = 0$, for all $i \geq 0$, then it is the unobservable model presented by Edelson and Hilderbrand [2]. For the rest of the paper, the former information disclosure policy is denoted u_+ and the latter policy is denoted u_-.

The provider's objective is to maximize its revenue, generated through the service completion. Thus, it aims to maximize the utilization of the queue which is equivalent to minimizing the idle stationary probability of the queue, denoted by π_0. The optimization problem for the provider is the following:

$$\min_{u \in \mathcal{U}} \pi_0^u,$$

where $\mathcal{U}$ is the set of information disclosure policies/mapping from N to the interval $[0, 1]$. Based on Naor's model, we know that an informed customer will join if the queue length is strictly lower than the threshold $L = \lfloor R\mu \rfloor$. An uninformed customer makes her decision based on her expected sojourn time, denoted by W^U. We assume that the uninformed customers are aware of the policy used by the provider (this information can be obtained by trials or via exogenous sources). Thus, an uninformed customer can evaluate her sojourn time given the system parameters and the provider policy.

Getting results for all type of information disclosure policies can be complicated. Therefore, in this paper, we restrict our analysis to deterministic threshold information disclosure policies, which make sense in applicative contexts. The principle of revealing the queue length when the system is empty and not revealing it when the system is full seems intuitive. Revealing information when the queue is small should increase the incoming rate of customers. At the opposite, when the system is overloaded, the provider may aim not to scare incoming customers and give them the information about the queue length.

4. Equilibrium Analysis of Threshold Policies

In this section, we consider threshold policies, denoted $u_D(.)$, for which the provider informs all customers about the queue length if the actual queue length is below or equal some threshold D and does not inform them otherwise, that is,

$$u_D(i) = \begin{cases} 1 & \text{if } i \leq D, \\ 0 & \text{if } i > D. \end{cases}$$

Since an uninformed customer knows that the queue length is greater than D, her expected sojourn time, denoted $W^U(D, q)$, depends on the threshold D and on the decision of the other uninformed customers, denoted by q, where $q \in [0, 1]$ represents the probability that an uninformed customer joins the queue (all customers are identical, hence we only consider symmetric strategies). We denote an equilibrium solution by q^*.

If $D \geq L - 1$, uninformed customers will never join the queue as their expected sojourn times, conditioned on the fact that the queue length is at least L, is necessarily higher than the reward R (recall that $L = \lfloor R\mu \rfloor$). Thus, the unique equilibrium is $q^* = 0$ and a threshold policy with a threshold $D \geq L-1$ is equivalent to the observable model. Henceforth, we only consider threshold policies with $D \in \{0, 1, .., L - 2\}$ (which also means that we only consider $R\mu \geq 2$). Thus, at equilibrium, all informed customers join, while all uninformed customers join with probability q.

The evolution of the queue length process forms a Markov chain with transition rate λ from state i to state $i + 1$ if $i \leq D$ and transition rate λq otherwise. The transition rate from i to $i - 1$ is μ for all $i > 0$. The Markov chain is illustrated in Figure 1. This Markov chain is a birth-death process and the stationary distribution is given by

$$\pi_i = \begin{cases} \pi_0 \left(\frac{\lambda}{\mu}\right)^i & \text{if } i \leq D + 1, \\ \pi_0 \frac{\lambda^i q^{i-D-1}}{\mu^i} & \text{if } i > D + 1, \end{cases} \tag{1}$$

where π_i is the stationary probability that the queue length is equal to i. The idle stationary probability (i.e., the probability that the queue is empty) as a function of the threshold D and the joining probability q is given by:

$$\begin{aligned} \pi_0(D, q) &= \left[\sum_{i=0}^{D+1} \left(\frac{\lambda}{\mu}\right)^i + \sum_{i=D+2}^{\infty} \frac{\lambda^i q^{i-D-1}}{\mu^i} \right]^{-1} \\ &= \frac{(\lambda - \mu)(\lambda q - \mu)}{\lambda^2(-1 + q)(\lambda/\mu)^D - \lambda q\mu + \mu^2}. \end{aligned} \tag{2}$$
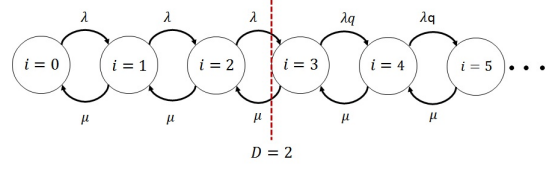

Figure 1: A Markov chain representation of the system evolution. D is the threshold above which customers are not informed about the queue length and q is the probability that uninformed customers join the queue. In this example, we assume $D < L - 1$, that is, informed customers always join the queue.

We define $\pi(i|i > D)$ to be the probability that the queue length is i, given that the queue length is higher than D. Using conditional probability rules, we obtain

$$\pi(i|i > D) = \frac{\pi_i}{\sum_{j=D+1}^{\infty} \pi_j}. \tag{3}$$

The expected sojourn time for an uninformed new customer is

$$W^U(D, q) = \sum_{i=D+1}^{\infty} \frac{i + 1}{\mu} \pi(i|i > D). \tag{4}$$

Using Equation (1) and Equation (3), we can express the expected sojourn time:

$$W^U(D, q) = \frac{\sum_{i=D+1}^{\infty} \pi_i \frac{i+1}{\mu}}{\sum_{i=D+1}^{\infty} \pi_i} = \frac{1 + D}{\mu} + \frac{1}{\mu - \lambda q}. \tag{5}$$

We next aim to determine the equilibrium strategy q^* which determines if an uninformed customer joins the queue or not. First, we prove the existence and uniqueness of the equilibrium.

Lemma 1. *A game with a threshold information policy has a unique equilibrium.*

Proof. Joining when the queue length is smaller than L and balking otherwise is a dominant strategy for the informed customers, hence unique. Given the behavior of the informed customers, we next analyze the behavior of the uninformed customers.

From Equation (5), we deduce that the expected payoff of an uninformed customer that joins (i.e., $W^U(D, q)$) is decreasing with q, while the expected payoff of balking R does not depend on q. Hence, the two payoff functions either intersect once or do not intersect. In the latter case, we will have a unique pure equilibrium where all uninformed customers join ($q^* = 1$) or all of them balk ($q^* = 0$); depending on which expected payoff is greater. In the former case, the intersection point forms

a mixed equilibrium. The fact that the expected payoff of joining is decreasing with q, while the expected payoff of balking does not depend on q means that the game has the "avoid-the-crowd" property which guarantees that the equilibrium is unique (see Chapter 1 in [10]). □

We start the equilibrium analysis with the simple case $\rho < 1$ and $R > 1/(\mu-\lambda)$. In this case, if policy u_- is used, then all customers join the queue because the expected sojourn time when all the customers join is $1/(\mu - \lambda)$. This outcome is obviously optimal for the provider and no other policy can outperform it. In any other case, an equilibrium where all customers join does not exist, namely at equilibrium $q < 1$. Henceforth, we focus on that scenario which implies that either the system is overloaded, that is $\rho > 1$, or it is underloaded with bounded reward, that is $\rho < 1$ and $R < 1/(\mu - \lambda)$.

The unique equilibrium might be a pure equilibrium with $q^* = 0$ or a mixed equilibrium with $q^* \in (0, 1)$. In the case of a pure equilibrium, an uninformed customer knows that the queue length is exactly $D + 1$ (since no uninformed customer joins the queue), and hence her expected sojourn time is $(D+2)/\mu$. In order for $q = 0$ to be an equilibrium, an uninformed customer should not be better off by joining, which implies that

$$D + 2 \geq R\mu \geq \lfloor R\mu \rfloor = L. \quad (6)$$

Since we only consider threshold policies with $D \leq L - 2$, we deduce that a pure equilibrium exists only if $R\mu = D + 2$.

Next, we study the case of having a mixed equilibrium. We derive the fraction of uninformed customers that join the queue, by using the property that at a mixed equilibrium, each player must be indifferent between the actions of joining and not joining the queue. Hence, at equilibrium, $W^U(D, q^*) = R$. By isolating q in Equation (5), we get

$$q^* = \frac{\mu(R\mu - D - 2)}{\lambda(R\mu - D - 1)}. \quad (7)$$

Since a pure equilibrium with $q^* = 0$ only exists when $R\mu = D+2$, the equation above captures both the case of a mixed equilibrium and the case of a pure equilibrium.

From Equation (7), we obtain the customers' equilibrium strategy under an information disclosure policy u_D adopted by the provider. This result is summarized with the following theorem.

Theorem 1. *If a provider uses the information disclosure policy $u_D(\cdot)$, all informed customers join the queue if the queue length is strictly smaller than L. Uninformed customers join the queue with probability $q^*(D)$, where*

$$q^*(D) = \begin{cases} 0, & \text{if } D \geq L - 1, \\ \frac{\mu(R\mu-D-2)}{\lambda(R\mu-D-1)}, & \text{otherwise.} \end{cases}$$

5. Comparison of Policies

Using Theorem 1, we next determine the optimal information disclosure policy that the provider should adopt in order to optimize its revenue. Toward this end, we derive and compare the idle stationary probabilities, at equilibrium, for the three types of policies: u_D, u_- and u_+. We ignore the case of $D \geq L-1$ since it is equivalent to the policy u_+.

For deriving the idle stationary probability when a threshold policy u_D is used, we substitute Equation (7) into Equation (2) (i.e., replacing q with q^*). We obtain that at equilibrium

$$\pi_0^*(D) = \frac{\mu - \lambda}{\mu - \lambda\rho^D[\rho(R\mu - 1 - D) - (R\mu - 2 - D)]}. \quad (8)$$

When sharing queue length information with all customers, we have, at equilibrium, an $M/M/1/L$ queue (i.e., a queue with finite capacity L). This type of queue was studied in [11] and the idle stationary probability is

$$\pi_0^{u_+} = \frac{\mu - \lambda}{\mu - \lambda\rho^{\lfloor R\mu \rfloor}}. \quad (9)$$

Finally, when no information is shared with customers there is a unique mixed equilibrium (recall that $R > 1/\mu$ and either $\rho > 1$ or $R < 1/(\mu - \lambda)$) and, hence, there is no pure equilibrium in this case). The effective arrival rate is λq^*. Since all customers are indifferent between joining and balking, the following holds:

$$R = \frac{1}{\mu - \lambda q^*}. \quad (10)$$

By isolating q^* and substituting it in the equation $\pi_0^{u_-} = 1 - \lambda q^*/\mu$, we get that

$$\pi_0^{u_-} = \frac{1}{R\mu}. \quad (11)$$

To compare between the performances of the different policies, we distinguish between the case of an overloaded queue (i.e., $\rho > 1$) and that of an underloaded queue (i.e., $\rho < 1$). Note that the strategic behavior of
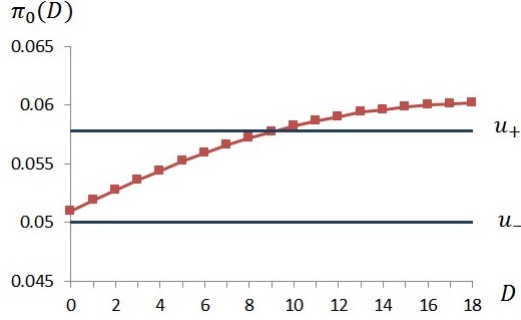
Figure 2: A queue with $\lambda = 9.8$, $\mu = 10$ and $R = 2$. The optimal policy is u_-.

Figure 3: A queue with $\lambda = 11$, $\mu = 10$ and $R = 2$. The optimal policy is u_+.

the customers guarantees that the effective load is always smaller than one.

Next, we show that in an overloaded queue, for any value of R and $D \in \{0, 1.., L - 2\}$, we always have $\pi_0^*(D) > \pi_0^{u_+}$. Similarly, in an underloaded queue case, $\pi_0^*(D) > \pi_0^{u_-}$ always holds. Hence:

Theorem 2. *Considering the set of deterministic threshold based information disclosure policies. To maximize its revenue, a service provider must either use the full information policy u_+ or the empty information policy u_-.*

Proof. See Section 6. □

In other words, it is never optimal to use a partial information policy, based on the queue length, which is somewhat counter-intuitive.

Figure 2 shows the stationary idle probability when different threshold values are used in an underloaded queue with $\lambda = 9.8$, $\mu = 10$ and $R = 2$. In this case $L = 20$. Thus, we consider threshold policies from 0 to 18. We can see from the figure that the idle stationary probability of any threshold policy is always greater than u_-. For some threshold values, it is also greater than u_+.

In Figure 3, we consider an overloaded queue with $\lambda = 11$, $\mu = 10$ and $R = 2$. In this case, the idle stationary probability of any threshold policy from 0 to 18 is bounded between u_- and u_+. This time, the policy that minimizes the idle probability is to always inform customers about the queue length, i.e. u_+.

6. Proof of Theorem 2

We split the proof into two cases. In the first case, we assume that $\rho > 1$. For this case, we will show that
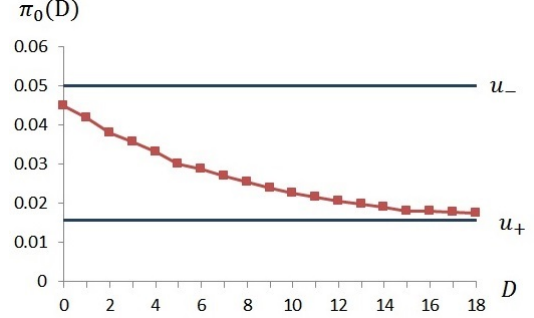
$\pi_0(D) > \pi_0^{u_+}$. Then, we study the case of $\rho < 1$, for which we show that $\pi_0(D) > \pi_0^{u_-}$.

From Equation (8) and Equation (9), we deduce that we can prove the first case by showing that:

$$\rho^{\lfloor R\mu \rfloor} > \rho^D[\rho(R\mu - 1 - D) - (R\mu - 2 - D)], \quad (12)$$

which is equivalent to showing that

$$\rho^{\lfloor R\mu \rfloor - D} > (\rho - 1)(R\mu - 1 - D) + 1. \quad (13)$$

We define $M = R\mu - D - 1$ and observe that

$$\rho^{\lfloor R\mu \rfloor - D} > \rho^{R\mu - D - 1} \quad (14)$$

(recall that $\rho > 1$). Thus, we can prove our claim by showing that for any $M \in [R\mu - \lfloor R\mu \rfloor + 1, R\mu - 1]$ and $\rho > 1$,

$$f(M) \triangleq \rho^M - (\rho - 1)M - 1 \geq 0. \quad (15)$$

We note that the smallest possible value of M is 1. Since $f(1) = 0$, showing that the derivative of $f(M)$ is non-negative will be the final step for proving the first case. The derivative of $f(M)$ with respect to M is given by

$$f' = \rho^M \ln(\rho) - \rho + 1. \quad (16)$$

Thus, we need to show that

$$\ln(\rho) \geq \frac{\rho - 1}{\rho^M}. \quad (17)$$

Using the First Mean Value Theorem for Integrals [12], one can show that $\ln(x) \geq (x - 1)/x$, for any $x > 0$. Thus, Equation (17) holds true and we conclude that the derivative of $f(M)$ is non-negative.

Next, we prove the other case where $\rho < 1$. We will show that

$$\frac{\mu - \lambda}{\mu - \lambda(\rho)^D[\rho(R\mu - 1 - D) - (R\mu - 2 - D)]} > \frac{1}{R\mu}, \tag{18}$$

which is equivalent to showing that

$$R(\mu - \lambda) > 1 - \rho^{D+1}[\rho(R\mu - 1 - D) - (R\mu - 2 - D)]. \tag{19}$$

We define $N = D + 1$ and with some algebra we get that Equation (19) is equivalent to

$$R(\mu - \lambda) > 1 + \rho^N[(1 - \rho)(R\mu - N) - 1]. \tag{20}$$

Next, we show that the equation above holds true for $N = 1$, which is the smallest possible value of N. We set $N = 1$ and we get:

$$R\mu(1 - \rho) > 1 + \rho[(1 - \rho)(R\mu - 1) - 1]. \tag{21}$$

With some algebra we get that the equation above is equivalent to:

$$R\mu - 1 > \rho(R\mu - 1). \tag{22}$$

The equation above holds true for any $R\mu \geq 2$ and for any $\rho < 1$. Hence, we showed that Equation (18) holds true for the case $N = 1$. Next, we will show that the right hand side of Equation (20) is non-increasing with N, which will conclude our proof. We define

$$g(N) = 1 + \rho^N[(1 - \rho)(R\mu - N) - 1]. \tag{23}$$

The derivative of $g(N)$ is

$$g' = -\rho^N \ln(\rho) + \rho^N \ln(\rho)(1 - \rho)(R\mu - N) - \rho^N(1 - \rho). \tag{24}$$

To show that the derivative is negative, we divide it by ρ^N and we need to show that:

$$-\ln(\rho) + \ln(\rho)(1 - \rho)(R\mu - N) - 1 + \rho \leq 0. \tag{25}$$

Since $\ln(\rho)(1 - \rho) < 0$ and $(R\mu - N) \geq 1$, the following holds:

$$\ln(\rho)(1 - \rho)(R\mu - N) \leq \ln(\rho)(1 - \rho). \tag{26}$$

Thus, we can show that Equation (25) holds true by showing that

$$-\ln(\rho) + \ln(\rho)(1 - \rho) - 1 + \rho \leq 0. \tag{27}$$

This is equivalent to showing that

$$\ln(\rho) \geq \frac{\rho - 1}{\rho}, \tag{28}$$

which holds for any $\rho > 0$, as we explained before.

QED

References

[1] P. Naor, The regulation of queue size by levying tolls, Econometrica: journal of the Econometric Society (1969) 15–24.

[2] N. M. Edelson, D. K. Hilderbrand, Congestion tolls for poisson queuing processes, Econometrica: Journal of the Econometric Society (1975) 81–92.

[3] O. Aksin, M. Armony, V. Mehrotra, The modern call-center: a multi-disciplinary perspective on operations management research, Production Operation Management 16 (2007) 665–688.

[4] R. Hassin, R. Roet-Green, Equilibrium in a two dimensional queueing game: When inspecting the queue is costly, Tech. rep., Working paper, Tel Aviv University, Israel (2011).

[5] R. Shone, V. A. Knight, J. E. Williams, Comparisons between observable and unobservable M/M/1 queues with respect to optimal customer behavior, European Journal of Operational Research 227 (1) (2013) 133–141.

[6] R. Hassin, Consumer information in markets with random product quality: The case of queues and balking, Econometrica: Journal of the Econometric Society (1986) 1185–1195.

[7] M. Hu, Y. Li, J. Wang, Efficient ignorance: Information heterogeneity in a queue, Rotman School of Management Working Paper (2472107).

[8] J. H. Large, T. W. Norman, Markov perfect bayesian equilibrium via ergodicity, paper in work.

[9] W. Stein, A. Rapoport, D. Seale, H. Zhang, R. Zwick, Batch queues with choice of arrivals: Equilibrium analysis and experimental study, Games and Economic Behavior 59 (2007) 345–363.

[10] R. Hassin, M. Haviv, To queue or not to queue: Equilibrium behavior in queueing systems, Vol. 59, Springer Science & Business Media, 2003.

[11] L. Takács, Introduction to the Theory of Queues, Vol. 584, Oxford University Press New York, 1962.

[12] P. R. Halmos, Measure theory, new york, 1950, Mathematical Reviews (MathSciNet): MR11: 504d.