



Queueing models for full-flexible multi-class call centers with real-time anticipated delays

Oualid Jouini, Yves Dallery, Zeynep Aksin

► To cite this version:

Oualid Jouini, Yves Dallery, Zeynep Aksin. Queueing models for full-flexible multi-class call centers with real-time anticipated delays. International Journal of Production Economics, 2009, 10.1016/j.ijpe.2008.01.011 . hal-01264963

HAL Id: hal-01264963

<https://hal.science/hal-01264963>

Submitted on 3 Feb 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Queueing Models for Full-Flexible Multi-class Call Centers with Real-Time Anticipated Delays

Oualid Jouini[†] Yves Dallery[†] Zeynep Aksin[‡]

[†] Ecole Centrale Paris Department of Industrial Engineering Grande Voie des Vignes 92290 Châtenay-Malabry, France walid.jouini@ecp.fr, yves.dallery@ecp.fr	[‡] Koç University College of Administrative Sciences and Economics Rumeli Feneri Yolu 34450 Sariyer-Istanbul, Turkey zaksin@ku.edu.tr
---	---

To appear in *International Journal of Production Economics*

October 2008

Abstract

In this paper, we consider two basic multi-class call center models, with and without reneging. Customer classes have different priorities. The content of different types of calls is assumed to be similar allowing their service times to be identical. We study the problem of announcing delays to customers upon their arrival. For the simplest model without reneging, we give a method to estimate virtual delays that is used within the announcement step. For the second model, we first build the call center model incorporating reneging. The model takes into account the change in customer behavior that may occur when delay information is communicated to them. In particular, it is assumed that customer reneging is replaced by balking that depends on the state of the system in this case. We develop a method based on Markov chains in order to estimate virtual delays of new arrivals for this model. Finally, some practical issues concerning delay announcement are discussed.

Keywords call centers, predicting delays, announcing delays, impatient customers, Markov chains, multiserver queues.

1 Introduction

Recent developments in technology and the business environment have dramatically increased the need for improved service systems. Well designed service systems allow to reduce costs and to

promote user satisfaction. The service sector is gaining prominence both in terms of the number of researchers and practitioners working in it, and its contribution to the gross domestic product, see Artiba (5). This paper deals with a well known type of service system, namely call centers. Call centers are used to provide services in many areas and industries: emergency centers, information centers, help-desks, tele-marketing and more. A telephone service enables customers to obtain a fast response, with a minimal effort. Providing services via call centers, instead of a face-to-face service, usually translates into lower operational costs to the service provider. The call center industry has been steadily growing and it has been observed worldwide. Estimations indicate that around 70% of all customer transactions occur in call centers, see Nakibly (22). Today, all Fortune 500 companies have at least one call center. They employ an average of 4,500 agents across their sites. More than \$300 billion is spent annually on call centers around the world, see McKinsey Quarterly (11).

The continued growth of call centers has brought with it a rich and interesting set of questions for both practitioners and academic researchers. This paper is motivated by a large French call center. As in many other organizations, the call center of our company constitutes the main point of contact with customers. Such centers have limited resources and face highly unpredictable demand that often result in long waits for customers. To improve customer satisfaction and alleviate congestion, call centers have recently started experimenting by informing arriving customers about anticipated delays, see Armony and Maglaras (3).

Information regarding anticipated delays is of a special importance in service systems with invisible queues, as in call centers. In such systems, the uncertainty involved in waiting is higher than that in visible queues, and it does not decrease over time. Customers have no means to estimate queue lengths or progress rate. So, feelings of frustration and anxiety increase during the wait. We expect that delay information would avoid such situations, and make the waiting experience more acceptable. Zakay (30) stipulates that waiting information may distract customers' attention from the passage of time. Hence, they may perceive the length of the wait as short.

The goal of this paper is to study announcing delays under different call centers situations. We consider two different multi-class call centers models, with and without abandonments. Since we are dealing with stochastic systems, there is no possible way to predict exact waiting times. The best one can do is to estimate their distribution. The service provider should then decide what the exact information that will be provided to customers is. For example, he may decide to provide the mean

of the estimated waiting time or any percentile of its distribution function. However, we should be careful: On the one hand, informing of a short waiting time, which is likely to underestimate the actual waiting, might reduce the reliability of the service provider in the eyes of the customers. On the other hand, informing of a long waiting time, might result in longer perceived waiting times and in a decrease in satisfaction.

For each call center configuration, we develop a method to estimate virtual delays of new arrivals. Since the goal is to provide information which is relevant to a specific type of customer at a specific time, we focus on estimating the waiting time given the system state at the time of estimation. This is different from estimating the overall performance of the system, such as the average waiting time of all customers, which is usually done assuming steady state. The models under consideration in this paper are formulated to be relevant in practice. The two major distinguishing features are priorities and the possibility for customers to be impatient. Priority mechanisms are a useful scheduling method that allows different customer types to receive differentiated performance levels. They are in addition known for their ease of implementation in practice. As for reneging, incorporating it in modeling is of great value in order to be as close as possible to reality. Waiting customers in call centers may naturally hang up once they feel that the time they spend in queue is too long, see Mandelbaum and Zeltyn (21).

The main contribution of this paper is the construction of a multi-class call center model with impatient customers that incorporates delay information. We propose a model in which the original behavior of customers (reneging) is substituted by balking upon arrival. To fully characterize the new model, we compute for each type of customer the balking probabilities and derive closed form expressions for the moments of their distribution of virtual delays. We note that the analysis here may be viewed as an extension of the work of Whitt (28). In the latter, the author addressed a similar problem for a call center model with a single class of impatient customers.

Here is how the rest of the paper is organized. In Section 2, we review some literature close to our work. In Section 3, we consider a multi-class call center model with infinitely patient customers. In Sections 3.1 and 3.2, we describe the model and develop a method for estimating virtual delays, respectively. This development would be used within the announcement step. We then move on and let customers renege while waiting in queue. In Section 4.1, we first describe the original model of the call center without delay information. In Section 4.2, we next focus on building the call center model assuming that it provides delay information to customers. In Section 4.3, we again develop

a method to derive virtual delays for each type of customer. In Section 5, some practical issues are discussed. In Section 5.1, we question the need of announcing delays when queues are empty upon the arrival of a new customer. In Section 5.2, we present a useful approximation of virtual delays in order to simplify the implementation of the announcement of delays. We conclude the paper in Section 6, where we highlight some directions for future research.

2 Literature Review

The continued growth of both importance and complexity of modern call centers has led to an extensive and growing literature. Due to the uncertainty governing the call center environment (customer and agent behaviors), the literature has typically addressed its issues using stochastic models, and in particular queueing models. Important related surveys are the paper of Koole and Mandelbaum (19) and its extended version Gans et al. (9) where the authors survey the literature dealing with the operations management of call centers.

The literature related to our work spans three main areas. The first is concerned with the analysis of multiserver systems motivated by call centers. The second deals with reneging phenomena. The third area is related to the problem of predicting and announcing delay information.

Let us focus on the first area, i.e., call center modeling. Call centers may be broadly classified into two contexts: multi-skill call centers and full-flexible call centers. A multi-skill call center handles several types of calls, and agents have different skills. The typical example, see Gans et al. (9), is an international call center where incoming calls are in different languages. Related studies include those by Akşin and Karaesmen (1), Chevalier and Tabordon (8), and references therein. Our concern in this paper is full-flexible call centers. In such a setting assistance to customers can be provided by any agent. This is a plausible assumption for many real cases, especially for unilingual call centers where the complete flexibility is not as difficult as in multilingual call centers. Furthermore, we assume for the models under consideration that all agents are statistically identical. In other words, they can answer all questions coming from customers with the same efficiency, both quantitatively and qualitatively, even in case of different types of customers. Our motivation is related to the nature of the call center we are considering here, and which is the case for many other call centers applications. The difference between customer types is only qualitative, i.e., it is not related to the statistical behavior of customers but to their importance for the company. Full-flexible call centers were extensively studied in the literature. We refer the reader to Gans et

al. (9), Jouini et al. (16), and references therein.

We now consider the second area of literature related to this paper, i.e., reneging phenomena. Queueing models incorporating impatient customers have received a lot of attention in the literature. They are an important feature in a wide variety of situations that may be encountered in manufacturing systems of perishable goods, telecommunication systems, call centers, etc. To underline the importance of abandonment modeling in the call center field, Gans et al. (9) and Mandelbaum and Zeltyn (21) give some numerical examples that point out the effect of abandonments on performance. The literature on queueing models with reneging focus especially on performance evaluation. We refer the reader to Ancker and Gafarian (2), Garnett et al. (10), and references therein for simple models assuming exponential reneging times. Garnett et al. (10) studies the subject of Markovian abandonments. They suggest an asymptotic analysis of their model under the heavy-traffic regime. Their main result is to characterize the relation between the number of agents, the offered load and system performance such as the probability of delay and the probability to abandon. Zohar et al. (31) investigates the relation between customer reneging and the experience of waiting in queue. Other papers have allowed reneging to follow a general distribution. Related studies include Baccelli and Hebuterne (6), Brandt and Brandt (7), Ward and Glynn (27), and references therein.

In what follows, we mention some literature related to the third area, namely predicting and announcing delays to customers. Predicting virtual delays in Markovian models, as under consideration in this paper, deals in particular with the transient analysis of birth-death processes. Several papers have been proposed for the study of the transient behavior of queues using birth-death processes, but in general, analytical solutions are extremely difficult to obtain. A number of interesting results are derived in Whitt (29). We also refer the reader to Jouini and Dallery (14), and Nakibly (22) for further results in various queueing situations. The next natural step afterwards is the announcement of predicted delays. We mention the relevant work of Whitt (28) where the author studies the effect of announcing delays on the performance of a single class call center. An extension of the latter work is addressed by Jouini et al. (15). Other close references include those by Armony et al. (4), Guo and Zipkin (12), and references therein. We also refer the reader to Vosz and Witt (25), and Kaminsky and Kaya (17) for work dealing with quoting customer lead times in the context of supply chains. In such studies, we should not ignore possible reactions by customers. Indeed, the announcement of delays would have a significant influence on

customers. The literature on customers influenced by delay information begins with Naor (23). An overview of customer psychology in waiting situations, including the impact of uncertainty, can be found in Maister (20). Taylor (24) showed that delays affect customers' service evaluations in an experiment involving airline flights. Hui and Tse (13) conducted a survey on the relationship between information and customer satisfaction.

3 Infinitely Patient Customers

In this section, we consider a call center model where customers are infinitely patient. The model is described in Section 3.1. Then in Section 3.2, we assume that the service provider gives delay information to customers upon arrival, and develop a method for estimating virtual delays. The latter would be used within the announcement step.

3.1 Model Description and Notation

Consider the queueing model of a call center with two classes of customers; valuable customers type A , and less valuable ones type B . The model consists of two infinite priority queues type A and B , and a set of s parallel, identical servers representing the set of agents. The call center is operated in such a way that at any time, any call can be addressed by any agent. So upon arrival, a call is addressed by one of the available agents, if any. If not, the call must join one of the queues. The scheduling policy of service assigns customers A (B) to queue A (B). Customers in queue A have priority over customers in queue B in the sense that agents are providing assistance to customers belonging to queue A first. The priority rule is non-preemptive, which simply means that an agent currently serving a customer pulled from queue B , while a new arrival joins queue A , will complete this service before turning to the queue A customer. Within each queue, customers are served in FCFS manner.

Arrival processes of type A and B customers follow a Poisson process with rates λ_A and λ_B , respectively. Let λ_T be the total arrival rate, $\lambda_T = \lambda_A + \lambda_B$. Successive service times are assumed to be i.i.d., and follow a common exponential distribution with rate μ for both types of customers. Then, the server utilization ρ (proportion of time each server is busy) is $\rho = \lambda_T/s\mu$. The condition for stability is $\rho < 1$, that is to say that the mean total arrival rate must be less than the mean maximal service rate of the system. The resulting model, referred to as Model 1, is shown in Figure 1.




Figure 1: Call center model without reneging, Model 1

There are two reasons for considering common distributions for service times. The first one relates to the types of call centers that motivate our analysis. We are considering call centers where customers are segmented into different groups based on their value to the firm. This segmentation can be based on lifetime value or profitability. The call center then provides different levels of service to these groups. This type of service level differentiation is widely used in financial service and telecommunication call centers. In the presence of this type of segmentation, the difference between customer types is not related to the statistical behavior of customers but to their importance for the company, which we capture through priorities. In concrete terms, we assume for our models that customer queries do not differ from one type of customer to another. The second reason is due to the complexity of the analysis when assuming different behaviors in the statistical sense. Our main purpose in this paper is to investigate simple but at the same time interesting models that allow us to better understand the system behavior and obtain practical guidelines.

In Section 3.2, we tackle the analysis of our call center by adding delay information. From the quantitative side, announcing delays would not affect our original model (Model 1), since customers are infinitely patient anyway. So, informing customers about their anticipated delays or not, leads to the same quantitative performance measures. However, announcing delays would ameliorate the waiting experience of customers by reducing uncertainty, and improve as a consequence their satisfaction.

3.2 Predicting and Announcing Virtual Delays

Consider a new arrival call. There are two possibilities: either at least one server is idle, or all servers are busy. In the former case, the customer enters service immediately without having to wait. So, the service provider does not announce any information to the customer. In the second case, he has to wait in queue for service to begin. In the following we give the distribution of the waiting time of a new arrival. This analysis will be used by the service provider afterwards, in

order to inform customers about their delays.

As it is the case in most modern call centers, we assume that the technology allows us to know the system state at each new arrival epoch. The system state at each new arrival epoch is defined by the number of customers in system. If the latter is larger than or equal to the number of servers s , then all servers are busy and the new arrival has to wait in queue. Let n_A be the number of type A customers in queue A seen by the new arrival, and n_B that of customers B in queue B , $n_A, n_B \geq 0$. Finally, let n_T be the total number of customers in queues seen by our new customer, $n_T = n_A + n_B$.

In what follows, we compute the mean and variance of the state-dependent waiting time distribution of each new customer type. At the epoch of each arrival, we assume that all servers are busy and $n_T = n_A + n_B \geq 0$ customers are waiting in queues. We separate the study depending on whether the call of interest is of type A or B . Type A customers observe a regular queue without priority, so estimation of their waiting time is easy to obtain. It is not the case for type B customers, because their waiting time is affected by future type A arrivals.

3.2.1 Virtual Delays for type A customers

Because of the strict priority, the waiting time of a new type A arrival does not depend on the number of type B customers waiting in queue. Also, it does not depend on future type A arrivals because queue A is working under the FCFS basis. Given n_A customers waiting in queue A , the call of interest has to wait until all n_A customers ahead of him enter service, plus the time it takes up to a service completion (of a customer in service). A service completion is exponentially distributed with mean $1/s\mu$ (s servers in parallel), and it is independent of the previous history because of the memoryless property. Hence, the waiting time of our customer is the sum of $n_A + 1$ i.i.d. exponential random variables each with mean $1/s\mu$, which has an Erlang distribution. Let us define $X_{n_A}^A$ as the conditional random variable measuring the waiting time of our customer, given the queue state n_A . The mean, $E(X_{n_A}^A)$, and variance, $Var(X_{n_A}^A)$, of $X_{n_A}^A$ are thereafter given by

$$E(X_{n_A}^A) = \frac{n_A + 1}{s\mu} \quad \text{and} \quad Var(X_{n_A}^A) = \frac{n_A + 1}{(s\mu)^2}. \quad (1)$$

We can also calculate the full probability distribution function (PDF) of $X_{n_A}^A$. The Erlang PDF is available in closed form, see for example Kleinrock, page 72 (18). Consider the ratio

$$\frac{\sigma(X_A)}{E(X_A)} = \frac{1}{\sqrt{n_A + 1}}, \quad (2)$$

which is independent of μ and s . We deduce that the waiting time distribution (conditional on all servers being busy) is highly concentrated about its mean, for large values of n_A . Note that the analysis above (for customers A) is still valid for the $GI/M/s$ queue. We only need to know the current state information.

3.2.2 Virtual Delays for type B customers

We focus on the waiting time distribution of a type B arrival who finds n_A type A customers and n_B type B customers waiting in queues. Before entering service, this customer has to wait for the queue to become empty of $n_T = n_A + n_B$ customers and of all future type A customers who arrive in between. Since all customers are statistically identical and the system is workconserving, we should observe that the waiting time of our customer does not depend on the order of service of the customers ahead of him. Hence, the duration of interest can be divided into two parts: The first is the busy period opened by the customer in service. The second part is the sum of $n_A + n_B$ busy periods, each one opened by one of the customer in queue. The busy period is the one of an $M/M/s$ queue with a Poisson arrival rate of λ_A . It is defined as the time from an arrival of a customer to the system with only one idle server until the first time one of the servers becomes idle.

Let us define $X_{n_T}^B$ as the conditional random variable measuring the waiting time of our customer, given queue state n_T . Knowing that the remaining service time of one customer is independent from the already finished work (exponential service times), then the distribution of the busy period opened by the customer in service is identically distributed as the busy period opened by one of the customers from the queue. Finally, the new type B customer has to wait for $n_A + n_B + 1$ i.i.d. busy periods of an $M/M/s$ queue with arrival rate λ_A . The probability density function (pdf) of the busy period of an $M/M/1$ queue can be found for example in Kleinrock, page 218 (18). Then, with a little thought it should be clear that the busy period pdf of interest here is obtained by only substituting the capacity of service μ (in the case of an $M/M/1$ queue) by $s\mu$ (in the case of an $M/M/s$ queue). To get a closed-form expression of the $X_{n_T}^B$ pdf, the mathematics becomes

complicated ($n_T + 1$ -fold convolution of the busy period pdf). Fortunately, the mean of $X_{n_T}^B$ is fairly simple to obtain, by summing the busy periods means up to $n_T + 1$. The same approach is still valid for the variance computation of $X_{n_T}^B$ using in addition the independence between the random variables of the busy periods durations. To get the first two moments of the busy period, we simply evaluate, respectively, the negative derivative and the positive second derivative at zero of its PDF Laplace transform in the time t . The mean, $E(X_{n_T}^B)$, and variance $\sigma(X_{n_T}^B)$, of $X_{n_T}^B$ are given by

$$E(X_{n_T}^B) = \frac{n_T + 1}{s\mu - \lambda_A} \quad \text{and} \quad Var(X_{n_T}^B) = (n_T + 1) \frac{s\mu + \lambda_A}{(s\mu - \lambda_A)^3}. \quad (3)$$

Equation (4) shows the ratio of the standard deviation, $Var(X_{n_T}^B) = \sqrt{Var(X_{n_T}^B)}$, by the mean,

$$\frac{\sigma(X_{n_T}^B)}{E(X_{n_T}^B)} = \sqrt{\frac{s\mu}{(n_T + 1)(s\mu - \lambda_A)}}. \quad (4)$$

Once again, we observe that the conditional waiting time distribution is highly concentrated about its mean, for large values of n_T .

It is not too difficult to extend the analysis of the conditional waiting time to the general case of an arbitrary number of customer classes. For example for the third priority class analysis, it is equivalent to aggregate the first two classes into one equivalent class. Thereafter, we use the same analysis as that conducted above for type B customers. This completes the study of announcing delays for call centers with infinitely patient customers.

4 Finitely Patient Customers

In this section, we consider a new major feature in our modeling. We let customers renege while waiting in queue. In Section 4.1, we first describe the original model of the call center without delay information. In Section 4.2, we next focus on building the call center model assuming that it provides delay information to customers. In Section 4.3, we finally derive for the latter model a number of performance measures related to queueing delays.

4.1 Model Description and Notation

We address the analysis of a call center with a single group of identical agents, serving two classes of impatient customers, high and low priority classes. The model is identical to that described in

Section 3.1, however in addition we allow customers to be impatient. After entering the queue, a customer will wait a random length of time for service to begin. If service has not begun by this time he will renege and is considered to be lost. Times before reneging for customers type A (B) are assumed to be i.i.d., and exponentially distributed with rate γ_A (γ_B). In practice, one would expect customers A to be less patient than customers B , $\gamma_A > \gamma_B$. We note that abandonments make our system unconditionally stable. The resulting model, referred to as Model 2, is shown in Figure 2.




Figure 2: Call center model with reneging, Model 2

4.2 Call Center Modeling with Announcement

Assume moving from the call center described in Section 4.1 to a call center with delay announcement. On the contrary to a call center with infinitely patient customers, there is a modeling complexity when we provide delay information to customers, due to possible changes in their behavior. In this section, we investigate the impact of announcing delays on the customer abandonment experience. When we inform a customer about his anticipated delay, he will decide from the beginning, either to hang up immediately because he estimates that his delay is too long, or to start waiting in queue. In the latter case, there are two further possibilities. The first is that customers never abandon thereafter. The second possibility is that the customer patience will change, i.e., customers may abandon even if they had chosen to start waiting. Customers would have a patience behavior different from that in the original system (without announcement), depending on the information we provide to them, since the announced time might influence their expectations or perceptions regarding waiting time. We refer the reader to Armony et al. (4) and Guo and Zipkin (12) for further details on the subject.

Several forms of delay information are possible. The best is that we would give to a new customer his actual delay, which cannot be known in advance because it is random. The most natural in practice is that the service provider gives a certain percentile β of the virtual delay

distribution to each new arrival. The virtual delay is the time it takes for a server to become free for the customer of interest. In other words, it is the time until all higher priority customers ahead of the arrival leave the queue plus the duration of a service completion. Whitt (28) has considered a similar problem for a single class call center. He proposed a model incorporating announcement by assuming that a new customer who finds all servers busy balks with a given probability. Once a customer elects to wait in queue, he would never abandon thereafter.

We assume that each new arrival comes with its own deadline of time patience, and paralleling to the model of Whitt (28), we stipulate that a new customer elects to join the queue with the probability that a server becomes free for him (his virtual waiting time) before he would renege. This is exact only if we assume that the customer acts as if the delay information was his actual delay, which is not the case. We do not let customers renege once they join the waiting line. This assumption may be reasonable for high values of β , since the estimation of the anticipated delay should be fairly accurate in that case. For lower β values our delay predictions will exhibit an upward bias. The possibility of renegeing after joining the queue is modeled in Jouini et al. (15), albeit in a setting with a single customer class. The results in that paper show the close and complex relationship between announcement precision β , balking and renegeing. It is shown numerically under what conditions announcing delays will be preferred to a system without announcement. We leave the extension of renegeing from queue subsequent to delay announcements in multi-class models to future research.

Whitt (28) considers the case of announcing the mean delay to customers. In our setting the distribution of virtual delays are summations of n_A (for customers A) or at least n_T (for customers B) random variables which we expect to converge to Normal distributions as system size grows. As a result, the mean will coincide with the median in busy systems, suggesting that an announcement with $\beta = 50\%$ will be quite close to the mean. This announcement precision is however quite low, implying that customers will experience delays longer than what has been announced to them in many cases. In practice one would expect such customers to renege, which is not captured by our model. Thus, for our results to be applicable, one should consider settings where announcement precision is high with β values that exceed 50%.

Assume that a new arrival finds n_A waiting type A customers in queue A , and n_B waiting type B customers in queue B . Note that implicitly we are focusing on new arrivals finding all servers busy. If the number seen by an arrival is less than s , then the new arrival never balks

and enters service immediately. Let us come back to a new arrival finding all servers busy. It should be clear that the probability of balking for a type A new arrival depends only on n_A (due to the priority rule), say $p_{bk}^A(n_A)$. However, the probability of balking for a new type B arrival depends on the couple (n_A, n_B) , say $p_{bk}^B(n_A, n_B)$. Furthermore, we should not fall in the confusion of only considering it as a function of $n_T = n_A + n_B$. Having different values of n_A and n_B , so that $n_T = n_A + n_B$ is held constant, would affect the virtual delay distribution of the customer of interest. The reason is that with delay information, the arrival rate of type A customers, seen by our new type B customer, depends on the state of queue A . As a consequence, not considering the couple (n_A, n_B) to compute the balking probability of that customer would lead to a wrong result.

Let $Y_{n_A}^A$ be the random variable measuring the state-dependent virtual delay for a new type A arrival finding n_A waiting customers ahead of him. Let $Y_{(n_A, n_B)}^B$ be the one for a new type B arrival finding n_A and n_B waiting customers ahead of him in queues A and B , respectively. Furthermore, let $G_{n_A}^A(t)$ and $G_{(n_A, n_B)}^B(t)$ for $t > 0$ be the PDF of $Y_{n_A}^A$ and $Y_{(n_A, n_B)}^B$, respectively. Then, the call center provides upon arrival the values

$$D_{n_A}^A = (G_{n_A}^A)^{-1}(\beta), \quad \text{and} \quad D_{(n_A, n_B)}^B = (G_{(n_A, n_B)}^B)^{-1}(\beta) \quad (5)$$

to type A and B customers, respectively. The balking probabilities are computed as follows. We denote by T_A (T_B) the random threshold patience for type A (B). The random variable T_A (T_B) has an exponential distribution with rate γ_A (γ_B). The probability for a new type A arrival to balk is thereafter

$$p_{bk}^A(n_A) = P(T_A < D_{n_A}^A). \quad (6)$$

The one for a new type B arrival is in turn given by

$$p_{bk}^B((n_A, n_B)) = P(T_B < D_{(n_A, n_B)}^B). \quad (7)$$

Using the PDF of the exponential distribution, the following holds

$$p_{bk}^A(n_A) = 1 - e^{-\gamma_A \cdot D_{n_A}^A}, \quad \text{and} \quad p_{bk}^B((n_A, n_B)) = 1 - e^{-\gamma_B \cdot D_{(n_A, n_B)}^B}. \quad (8)$$

The resulting model of the call center, incorporating delay information, is shown in Figure




Figure 3: Call center model with delay information, Model 3

3, and is referred to as Model 3. Note that it is reasonable to assume that balking decisions are independent from one customer to another, so that arrivals still follow a modulated Poisson process. What remains to be done in order to fully characterize Model 3 is to compute state-dependent arrival rates for each customer type, which in turn reduces to characterizing the distribution functions of $Y_{n_A}^A$ and $Y_{(n_A, n_B)}^B$. In the next section, we give closed-form expressions for their first two moments. Based on these results, we propose in Section 5.2, a helpful and practical approximation of their entire distributions.

4.3 Predicting and Announcing Virtual Delays

As in Section 3.2, we assume that the technology of our call center enables us to know when queues are empty, and whether there is an available agent for an upcoming customer, or not. If less than s customers are present in the system, the customer of interest gets service immediately. If not, he has to wait in his corresponding queue for service to begin.

Knowing that all servers are busy, we focus on analyzing the conditional random variables $Y_{n_A}^A$ and $Y_{(n_A, n_B)}^B$, $n_A, n_B \geq 0$. We separate the analysis depending on whether the arrival call is of type A or B . Similarly to Section 3.2, the priority schemes under consideration makes the analysis for type A customers less complicated than that for type B customers. The latter is indeed affected by future type A arrivals who have higher priority for service.

Let us recall that we are calculating virtual delays which will be used within a second step in order to compute balking probabilities. In other words, we are calculating the time it takes until a server becomes free for the customer of interest in case he elects to wait (does not balk). In what follows, we analyze these quantities for both customer types.

4.3.1 Virtual Delays for Type A Customers

Consider a new type A arrival who finds all servers busy, n_A waiting customers in queue A and n_B waiting customers in queue B . Owing to his higher priority, the virtual delay of a new type A arrival does not depend on the number of type B customers already present in system, see Figure 4. The customer has to wait until the n_A waiting customers leave the queue plus the time it takes for a service completion (when all servers are busy). By a customer who leaves the queue, we only mean a customer who enters service. In Model 3, there is no longer the possibility for customers to renege once they elect to join the waiting line.




Figure 4: Virtual delay for a new type A arrival

Our customer A has to wait for the n_A customers ahead of him to enter service, and then he has to wait for a service completion. Overall, he has to wait for $n_A + 1$ service completions. Hence, the pdf of $Y_{n_A}^A$ is simply the convolution of the pdfs of $n_A + 1$ i.i.d. exponential random variables each with parameter $s\mu$. So, $Y_{n_A}^A$ has an $n_A + 1$ -Erlang distribution with parameter $s\mu$. The mean and variance of $Y_{n_A}^A$ are, respectively, given by

$$E(Y_{n_A}^A) = \frac{n_A + 1}{s\mu}, \quad \text{and} \quad \text{Var}(Y_{n_A}^A) = \frac{n_A + 1}{s^2\mu^2}. \quad (9)$$

Using the PDF of $Y_{n_A}^A$, it only remains to come back to Equations (5) and (6) in order to compute the balking probability $p_{bk}^A(n_A)$. Define now the standard deviation of $Y_{n_A}^A$ by $\sigma(Y_{n_A}^A) = \sqrt{\text{Var}(Y_{n_A}^A)}$, and the coefficient of variation by the ratio of the standard deviation over the mean, $cv(Y_{n_A}^A) = \sigma(Y_{n_A}^A)/E(Y_{n_A}^A)$. As shown in Equation (10), the ratio $cv(Y_{n_A}^A)$ is characterized to have a simple form independent of μ and s .

$$cv(Y_{n_A}^A) = \frac{1}{\sqrt{n_A + 1}}. \quad (10)$$

From Equation (10), we again note that for large values of n_A , the virtual delay of $Y_{n_A}^A$ is very concentrated about its mean. This implies that for large values of n_A , the mean value of $Y_{n_A}^A$ should

provide a good approximation of the virtual delay.

4.3.2 Virtual Delays for Type B Customers

Knowing that all servers are busy, let n_A and n_B be the number of type A and B waiting customers seen by a new type B arrival, in queues A and B , respectively.

The random variable $Y_{(n_A, n_B)}^B$ is the time until the $n_T = n_A + n_B$ waiting customers start service, plus the time it takes for all future type A arrivals (during the wait of the customer of interest) to enter service, plus the duration for a service completion (when all servers are busy), see Figure 5.

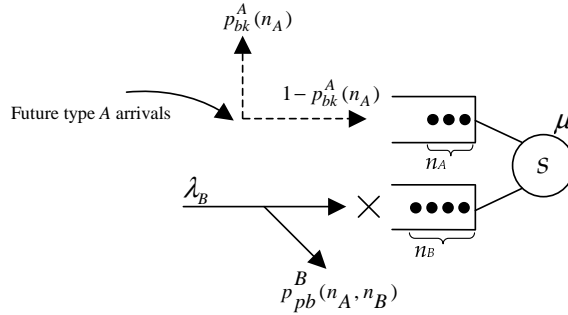


Figure 5: Virtual delay for a new type B arrival

To characterize $Y_{(n_A, n_B)}^B$, we ignore all future type B arrivals because the discipline of service within queue B is FCFS. However, all future type A arrivals have to be considered because of their higher priority against the customer of interest. Recall that reneging is no longer possible. We only consider events of type A arrivals and service completions. Thereby, changes of queue states seen by our customer are as follows. As long as type A customers are waiting in queue, the number of type B waiting customers does not change, however that of type A customers increases by one further to a type A arrival or decreases by one further to a service completion. The number of type B waiting customers cannot increase. It only decreases by one further to a service completion when no type A customers are waiting in queue. We should be careful not to forget that type A arrivals are state-dependent due to the balking decisions of customers upon arrival.

Based on the above explanation, we move on to employ the following two-dimensional Markov chain. Let the system state at a given random instant be (m_A, m_B) where m_A (m_B) is the number of type A (B) customers in queue A (B), $m_A, m_B \geq 0$. In addition, the Markov chain has an absorbing state denoted by (-1) . The system moves to (-1) subsequent to a service completion

when both queues are empty. Reaching the latter state means that a server is available for the customer of interest. When m_A customers are waiting in queue A , we denote the state-dependent arrival rates of type A arrivals by

$$\lambda_A(m_A) = \lambda_A(1 - p_{bk}^A(m_A)), \quad m_A \geq 0. \quad (11)$$

The non-zero transition rates are

$$\begin{cases} q(m_A, m_B)(m_A+1, m_B) = \lambda_A(m_A), \text{ for } m_A, m_B \geq 0, \\ q(m_A, m_B)(m_A-1, m_B) = s\mu, \text{ for } m_A, m_B > 0, \\ q(0, m_B)(0, m_B-1) = s\mu, \text{ for } m_B \geq 0, \\ q(0, 0)(-1) = s\mu. \end{cases} \quad (12)$$

As shown in Figure 6, measuring $Y_{(n_A, n_B)}^B$ may be formulated as the downcrossing time until absorption in state (-1) , starting from state (n_A, n_B) .

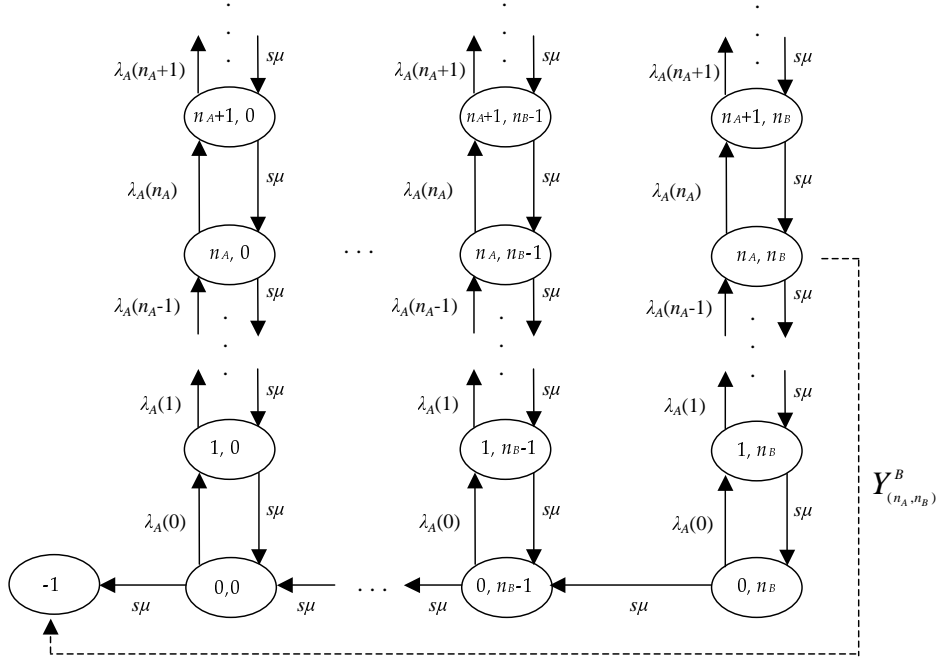


Figure 6: The random variable $Y_{(n_A, n_B)}^B$

The Markov chain we consider has a special structure allowing analytical solutions. From Figure 6, the random variable $Y_{(n_A, n_B)}^B$ may be rewritten as

$$Y_{(n_A, n_B)}^B = U(n_A) + V_{n_B-1} + \dots + V_0, \quad (13)$$

where $U(n_A)$ is the random variable measuring the downcrossing time until first passage at state $(0, n_B - 1)$ starting from state (n_A, n_B) , V_i is the random variable measuring the downcrossing time until first passage time at state $(0, i - 1)$ starting from state $(0, i)$ for $1 \leq i \leq n_B - 1$, and V_0 is the random variable measuring the downcrossing time until absorption in state (-1) starting from state $(0, 0)$.

The Markovian assumptions allow us to state that the random variables $U(n_A)$, V_0 , ..., and V_{n_B-1} are independent. From Figure 6, we see that V_0 , ..., and V_{n_B-1} are identically distributed. Let $E(Y_{(n_A, n_B)}^B)$ and $Var(Y_{(n_A, n_B)}^B)$ be the mean and variance of the random variable $Y_{(n_A, n_B)}^B$, respectively. Then, using the linearity property of expectations, we get

$$E(Y_{(n_A, n_B)}^B) = E(U(n_A)) + n_B E(V_0), \quad (14)$$

and from the independence between these random variables, the following holds

$$Var(Y_{(n_A, n_B)}^B) = Var(U(n_A)) + n_B Var(V_0). \quad (15)$$

Let us now focus on computing the means and variances of $U(n_A)$, V_0 , ..., and V_{n_B-1} . To do so, we define an intermediate birth-death process with discrete state space. The transition rates of the process are denoted by

$$\begin{cases} q_{m, m+1} = \lambda_A(m), \text{ for } m \geq 0, \\ q_{m, m-1} = s\mu, \text{ for } m \geq 1, \\ q_{0, a} = s\mu, \end{cases} \quad (16)$$

and $q_{m, n} = 0$ otherwise. The birth-death process is derived from the previous Markov chain and is shown in Figure 7.




Figure 7: Intermediate birth-death process

One may intuitively see that the intermediate birth-death process allows us to compute the time it takes to empty the queue of a given number of waiting type A customers plus the time for

a service completion, so that one of the waiting type B customer could enter service. The random variables $U(n_A)$, V_0 , ..., and V_{n_B-1} are defined on the intermediate birth-death process as follows. The random variable $U(n_A)$ is the downcrossing time until first passage at state a , starting from state n_A . As for the random variable V_i , $0 \leq i \leq n_B - 1$, it is only the first passage time at state a , starting from state 0.

By considering a general birth-death process, Jouini and Dallery (14) gives closed-form expressions for any moment of order $k \geq 1$ of several random variables related to first passage times. We use their results in our context here. To simplify the presentation, we introduce the quantities δ_m for $m \geq 0$. For $m = 0$ we let $\delta_0 = 1$ (one per unit of time, u.t.), and for $m \geq 1$ we let $\delta_m = \lambda_A(m - 1)$. Let us now define the potential coefficients of the intermediate birth-death process, say ϕ_m , as follows.

$$\phi_0 = 1, \text{ and } \phi_m = \frac{\prod_{j=0}^{m-1} \delta_j}{s^m \mu^m}, \text{ for } m \geq 1. \quad (17)$$

From Jouini and Dallery (14), the mean $E(U(n_A))$ and variance $Var(U(n_A))$ of $U(n_A)$ are then given by

$$E(U(n_A)) = \sum_{m=0}^{n_A} \frac{1}{\delta_m \phi_m} \sum_{j=m+1}^{\infty} \phi_j, \quad (18)$$

$$Var(U(n_A)) = \sum_{m=0}^{n_A} \frac{2}{\delta_m \phi_m} \sum_{j=m+2}^{\infty} \frac{1}{\delta_{j-1} \phi_{j-1}} \left(\sum_{l=j}^{\infty} \phi_l \right)^2 + \sum_{m=0}^{n_A} \frac{1}{\delta_m^2 \phi_m^2} \left(\sum_{j=m+1}^{\infty} \phi_j \right)^2. \quad (19)$$

The mean $E(V_0)$ and variance $Var(V_0)$ of the random variable V_0 are given by

$$E(V_0) = \frac{1}{\delta_0} \sum_{m=1}^{\infty} \phi_m, \quad (20)$$

$$Var(V_0) = \frac{2}{\delta_0} \sum_{m=2}^{\infty} \frac{1}{\delta_{m-1} \phi_{m-1}} \left(\sum_{j=m}^{\infty} \phi_j \right)^2 + \frac{1}{\delta_0^2} \left(\sum_{m=1}^{\infty} \phi_m \right)^2. \quad (21)$$

Note that we kept the intermediate parameter $\delta_0 (=1)$ in Equations (20) and (21) in order to get the correct unit of the quantities $E(V_0)$ and $Var(V_0)$. Substituting Equations (18), (19), (20) and (21) back into Equations (14) and (15) leads to the expressions of the mean and variance of the random variable $Y_{(n_A, n_B)}^B$. Finally, using the results of Section 4.3.1 to compute the balking probabilities for type A arrivals, the mean and variance of $Y_{(n_A, n_B)}^B$ are thereafter fully characterized.

Note that one may derive all higher order moments of the virtual delay for both customer types, which allows us to derive their full distributions. However, the analysis would be cumbersome and numerically time consuming. We thereafter content ourself with only the first two moments, and propose a useful approximation of these distributions as we shall explain later in Section 5.2. Furthermore, for finite buffer systems the same analysis holds and the infinite summations in the expressions for the first two moments are replaced by finite ones.

5 Some Practical Issues

In this section, we investigate some practical issues for an eventual implementation of delay information. We identify two points that may help practitioners. The first is discussed in Section 5.1, and evaluates the need of announcing delays when queues are empty upon the arrival of a new customer. The second is discussed in Section 5.2, and deals with an approximation for computing the anticipated delay we communicate to each new arrival.

5.1 Empty Queues

In this section, we call into question the need for communicating delays to new customers who find the queues empty. We consider a new arriving call finding empty queues. We define q as the conditional probability that the call of interest has to wait before beginning service, given that he finds the queues empty. The probability q does not depend on the type of the new arrival, and it concerns the case of having all servers busy knowing that the queues are empty. We are interested in calculating q , in order to get some indication on its value under normal working conditions of a call center. For instance, if the proportion q is very small, then most arriving calls begin service without waiting. So, there is little need for delay information. However, if q is quite large, then a considerable proportion of new calls have to wait, and the prediction would be important.

Another reason that prompts us to address this problem is related to an organizational issue in some call centers. In most call centers the state of the queues, i.e., the number of waiting customers if any, is visible. If no customers are waiting in queue, the system may however be unable to tell whether an agent is available for the new arrival or not. Even if technological capability to monitor agent availability is in place, agents are known to game such systems in practice. As a consequence, it would be interesting to study how the quantity q behaves.

The proportion q represents the conditional probability that a new call has to wait knowing

that the queues are empty. To compute q , the analysis is straightforward. We choose to omit the details. For the call center incorporating announcement of delays of Section 3 (Model 1, without reneging) and that of Section 4 (Model 3, with reneging), we have the following simple and common expression of q for both models

$$q = \frac{\lambda_T^s}{s! \mu^s} \left(\sum_{i=0}^s \frac{\lambda_T^i}{i! \mu^i} \right)^{-1}. \quad (22)$$

We have conducted a numerical analysis which we do not report here for brevity. In it, we consider a large range of parameters typical of those that would be encountered in real situations. The central insight is that only a small proportion of new calls (who find empty queues) have to wait before beginning service. Under normal working conditions, the proportion is of the order of 10%. As a consequence, there is not a real need to inform customers about their anticipated delays. In addition, even if one customer has to wait, his waiting time is sufficiently low.

5.2 Normal Approximation of Virtual Delays

Given the system state upon each arrival and given a coverage probability β , the service provider has to compute the value of the anticipated delay. This numerical computation operation is characterized to be too heavy. In fact for the low priority customers, exact moment expressions of the virtual delays involve infinite summations. Even if these are replaced by finite summations for the case with buffers, some numerical difficulties will remain since we need to conduct such real-time operations for each arrival. Off-line tables would have to be built to circumvent these problems.

From a practical point of view, a normal distribution provides a satisfactory approximation of virtual delays. Since the random variables of virtual delays we consider here deal with summations of independent random variables, the Normal approximation should work well, see Whitt (29) and Ward and Whitt (26). This claim is supported by theoretical results based on the Law of Large Numbers and the Central Limit Theorem. The Normal approximation should especially work well for new arrivals who find a large number of waiting customers in queue.

We only need the mean and standard deviation of the state-dependent virtual delay in order to get its full distribution (approximately). Thus, we propose to use the Normal distribution by only picking up the means and variances we derived in Sections 3 and 4. We should however point out that for small values of β , such distributions may lead to negative values of anticipated delays. To be judicious, we may adapt a given normal distribution by truncating it. For instance, let $h(t)$

and $H(t)$, $-\infty < t < +\infty$, be the pdf and PDF of the original Normal distribution, respectively. Also, let $h_{tr}(t)$ and $H_{tr}(t)$, $0 < t < +\infty$, be those for the truncated Normal distribution. The pdf of $h_{tr}(t)$ is calculated as $h_{tr}(t) = \frac{h(t)}{1-H(0)}$. By doing so, we even out the area of the negative region ($t < 0$) over that of the positive region, so that we build an appropriate distribution. Note that this transformation should not really affect the original Normal distribution. The reason is that the quantity $1 - H(0)$ is low for very small numbers of waiting customers in queue, and may be reasonably neglected otherwise.

To assess the quality of the normal approximation, we conduct a simple numerical analysis. We only consider the case of the more important type A customers in Model 3 (with reneging). We believe that the same conclusions remain valid for all remaining cases. We quantitatively compare the exact distributions of virtual delays of new arrivals A with the approximated ones (Normal distribution). As we have investigated in Section 4.3.1, the virtual delay of a new arrival A who finds n_A waiting customers A in queue, $Y_{n_A}^A$, has an $n_A + 1$ -Erlang distribution with parameter $s\mu$. The key parameters are thereafter n_A and $s\mu$. We choose $s\mu = 3$ (for example 15 agents and an average service time of 5 min). In Figure 8, we plot the pdf of the exact virtual delay and that of the approximated one for $n_A=1, 5$, and 10. From Figure 8, we deduce that the quality of the approximation improves as the number of waiting customers in queue increases. This agrees with our claim above.




Figure 8: Exact and approximated density functions

To assess the quality of the Normal approximation for a given percentile β , we plot in Figure 9 the PDF of the exact and approximated virtual delays. We choose $s\mu = 3$ and $n_A = 5$. From Figure 9, we see that the approximation is very accurate, especially for high values of confidence intervals. The manager will announce a delay which is based on a given β . Naturally, he should




Figure 9: Exact and approximated probability functions

preserve his credibility by choosing a high confidence interval (in the order of 90%). For example for $\beta = 90\%$, the approximated delay is 3.05 min while the exact one is 3.1 min.

6 Concluding Remarks and Further Extensions

We formulated and analyzed a multi-class call center model with priorities and impatient customers, to whom anticipated delays may be announced upon arrival. Assuming that delay announcements will lead customers to balk in response, the analysis analytically characterizes virtual delays that will be experienced in such call centers. While the results in the paper enable an exact analysis, an approximation based on a Normal distribution is also proposed, in case computational ease is valued more than announcement precision in practice. It is shown that especially for busy systems such an approximation will not sacrifice a lot of precision.

Worrying about accurate and practical results, one may continue our work by empirically describing customer reactions in response to delay announcements, in order to validate our modeling of that reaction herein. Guided by such results, one could then explore whether delays should be announced to customers, to which types of customers such announcements should be made, and with what level of precision. A further topic for future research is to analyze more complex systems: more than two customer classes and other disciplines of service such as threshold policies. It would also be of value to extend our work to the case of different statistical behaviors of different customer types. In other words, to settings where service times as well as times before reneging are not identically distributed for different customer types.

Acknowledgements

The authors were supported by The Scientific & Technological Research Council of Turkey, TÜBİTAK. The first author benefited from a scholarship provided by TÜBİTAK. The paper was written in part while the first author was visiting Koç University. The authors wish to thank the anonymous referees for their helpful suggestions to improve this paper.

References

1. O. Z. Akşın and F. Karaesmen. Designing Flexibility: Characterizing the value of Cross-Training Practices. 2002. Working paper, Koç University.
2. C. J. Ancker and A. Gafarian. Queueing with Impatient Customers Who Leave at Random. *Journal of Industrial Engineering*, 13:84–90, 1962.
3. M. Armony and C. Maglaras. Contact Centers with a Call-Back Option and Real-Time Delay Information. *Operations Research*, 52:527–545, 2004.
4. M. Armony, N. Shimkin, and W. Whitt. The Impact of Delay Announcements in Many-Server Queues with Abandonment. *Operations Research*, 2008. To appear.
5. A. Artiba. Performance Evaluation, Planning and Control, of Productive Systems. *International Journal of Production Economics*, 85:1–2, 2003.
6. F. Baccelli and G. Hebuterne. On Queues With Impatient Customers. *Performance'81 North-Holland Publishing Company*, pages 159–179, 1981.
7. A. Brandt and M. Brandt. Asymptotic Results and a Markovian Approximation for the $M(n)/M(n)/C + GI$ System. *Queueing Systems*, 41:73–94, 2002.
8. P. Chevalier and N. Tabordon. Overflow Analysis and Cross-Trained Servers. *International Journal of Production Economics*, 85:47–60, 2003.
9. N. Gans, G. Koole, and A. Mandelbaum. Telephone Call Centers: Tutorial, Review, and Research Prospects. *Manufacturing & Service Operations Management*, 5:73–141, 2003.
10. O. Garnett, A. Mandelbaum, and M. Reiman. Designing a Call Center with Impatient Customers. *Manufacturing & Service Operations Management*, 4:208–227, 2002.

11. K.A. Gilson and D.K. Khandelwal. Getting More from call Centers. 2005. McKinsey Quarterly: The Online Journal of McKinsey & Co. Available at <http://www.marketingpower.com/content25334.php>.
12. P. Guo and P. Zipkin. Analysis and Comparison of Queues with Different Levels of Delay Information. *Management Science*, 53:962–970, 2007.
13. M. Hui and D. Tse. What to Tell Customer in Waits of Different Lengths: an Integrative Model of Service Evaluation. *Journal of Marketing*, 60:81–90, 1996.
14. O. Jouini and Y. Dallery. Moments of First Passage Times in General Birth-Death Processes. *Mathematical Methods of Operational Research*, 68:49–76, 2007.
15. O. Jouini, Y. Dallery, and O. Z. Aksin. Call Centers with Delay Information: Models and Insights. 2008. Working paper. Ecole Centrale Paris.
16. O. Jouini, Y. Dallery, and R. Nait-Abdallah. Analysis of the Impact of Team-Based Organizations in Call Centers Management. *Management Science*, 54:400–414, 2008.
17. P. Kaminsky and O. Kaya. Inventory positioning, scheduling and lead-time quotation in supply chains. *International Journal of Production Economics*, 114:276–293, 2008.
18. L. Kleinrock. *Queueing Systems, Theory*, volume I. A Wiley-Interscience Publication, 1975.
19. G. Koole and A. Mandelbaum. Queueing Models of Call Centers An Introduction. *Annals of Operations Research*, 113:41–59, 2002.
20. D. Maister. Psychology of Waiting Lines. *Harvard Business School Cases*, pages 71–78, 1984.
21. A. Mandelbaum and S. Zeltyn. Staffing Many-Server Queues with Impatient Customers: Constraint Satisfaction in Call Centers. 2008. Working paper, Technion, Haifa, Israel.
22. E. Nakibly. *Predicting Waiting Times in Telephone Service Systems*. 2002. Ph.D. Thesis, The Senate of the Technion.
23. P. Naor. The Regulation of Queue Size by Levying Tolls. *Econometrica*, 37:15–24, 1969.
24. S. Taylor. Waiting for Service: The Relationship Between Delays and Evaluations of Service. *Journal of Marketing*, 58:56–69, 1994.

25. S. Vosz and A. Witt. Hybrid flow shop scheduling as a multi-mode multi-project scheduling problem with batching requirements: A real-world application. *International Journal of Production Economics*, 105:445–458, 2007.
26. A. R. Ward and W. Whitt. Predicting Response Times in Processor-Sharing Queues. *Proceedings of the Fields Institute Conference on Communication Networks*, 2000.
27. A.R. Ward and P.W. Glynn. A Diffusion Approximation for a Markovian Queue with Reneging. *Queueing Systems*, 43:103–128, 2003.
28. W. Whitt. Improving Service by Informing Customers about Anticipated Delays. *Management Science*, 45:192–207, 1999.
29. W. Whitt. Predicting Queueing Delays. *Management Science*, 45:870–888, 1999.
30. D. Zakay. An Integrated Model of Time Estimation. *Times and Human Cognition: A Life Span Perspective*, 1989. Iris Levin and Dan Zakay, eds, Amsterdam: North Holland.
31. E. Zohar, A. Mandelbaum, and N. Shimkin. Adaptive Behavior of Impatient Customers in Tele-Queues: Theory and Empirical Support. *Management Science*, 48:566–583, 2002.