

Statistical Analysis with Little's Law

Song-Hee Kim and Ward Whitt

Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027-6699,
{sk3116,ww2040}@columbia.edu, <http://www.columbia.edu/~{sk3116,ww2040}>

The theory supporting Little's law ($L = \lambda W$) is now well developed, applying to both limits of averages and expected values of stationary distributions, but applications of Little's law with actual system data involve measurements over a finite time interval, which are neither of these. We advocate taking a statistical approach with such measurements. We investigate how estimates of L and λ can be used to estimate W when the waiting times are not observed. We advocate estimating confidence intervals. Given a single sample path segment, we suggest estimating confidence intervals using the method of batch means, as is often done in stochastic simulation output analysis. We show how to estimate and remove bias due to interval edge effects when the system does not begin and end empty. We illustrate the methods with data from a call center and simulation experiments.

Key words: Little's law; $L = \lambda W$; measurements; parameter estimation; confidence intervals; bias; finite-time versions of Little's law; confidence intervals with $L = \lambda W$; edge effects in $L = \lambda W$; performance analysis.

History: submitted March 25, 2012, Revisions August 17 and December 20, 2012

1. Introduction

We have just celebrated the 50th anniversary of the famous paper by Little (1961) on the fundamental queueing relation $L = \lambda W$ with a retrospective by Little (2011) himself, which emphasizes the applied relevance as well as reviewing the advances in theory, including the sample-path proof by Stidham (1974) and the extension to $H = \lambda G$. Several books provide thorough treatments of the theory, including the sample-path analysis by El-Taha and Stidham (1999) and the stationary framework involving the Palm transformation by Sigman (1995) and Baccelli and Bremaud (2003), as well as the perspective within stochastic networks by Serfozo (1999). As a consequence, $L = \lambda W$ and the related conservation laws are now on a solid mathematical foundation.

The relation $L = \lambda W$ can be quickly stated: The average number of customers waiting in line

(or items in a system), L , is equal to the arrival rate (or throughput) λ multiplied by the average waiting time (time spent in system) per customer, W . If we know any two of these quantities, then we necessarily know all three. The easily understood reason is reviewed in §2. With queueing models where λ is known, the relation $L = \lambda W$ yields the value of L or W whenever the other has been calculated.

1.1. Measurements Over a Finite Time Interval

However, in many applications, these conservation laws are applied with measurements over a finite time interval of length t , yielding finite averages $\bar{L}(t)$, $\bar{\lambda}(t)$ and $\bar{W}(t)$ (defined in (1) below). Indeed, the applied relevance with measurements motivated Little (2011) to discuss relations among finite-time measurements instead of the stationary framework in Little (1961). However, with finite averages, the large body of supporting theory often does not apply directly, because that theory concerns either long-run averages (limits) or the expected values of stationary stochastic processes in stochastic models. The available measurements are neither of these.

Here is the essence of a typical application: We start with the observation of $L(s)$, the number of items in the system at time s , for $0 \leq s \leq t$. From that sample path, we can directly observe the arrivals (jumps up) and departures (jumps down). Hence, we can easily estimate the arrival rate λ and the average number in system L . However, based only on the available information, we typically cannot determine the time each item spends in the system, because the items need not depart in the same order that they arrived. Nevertheless, we can estimate the average waiting time by $W = L/\lambda$, using our estimates of L and λ .

In this paper we focus on the typical application in the paragraph above, estimating W given estimates of L and λ , illustrated by data from a large call center. The first issue is that, with commonly accepted definitions (see (1) below), the relation $L = \lambda W$ is not valid as an equality over a finite time interval unless the system starts and ends empty, which often is either not feasible or not desirable. In §2 we review the exact relation that holds for finite time intervals and a way to modify the definitions so that the edge effects do not occur, even when the system does not start and

end empty. Using modified definitions to make $\bar{L}(t) = \bar{\lambda}(t)\bar{W}(t)$ valid for *all* finite intervals is the approach of the “operational analysis” proposed by Buzen (1976) and Denning and Buzen (1978), motivated by performance analysis of computer systems, which is also discussed by Little (2011). Changing definitions in that way can be very helpful to check the consistency of measurements and data analysis, which is a legitimate concern. While changing the definitions is one option, we advocate *not* doing so, because it leads to problems with interpretation.

1.2. A Statistical Approach

We advocate taking a statistical approach with data over a finite time interval. Thus we regard the finite averages as realizations of random estimators of underlying unknown “true” values. We suggest estimating confidence intervals. Since the initial estimators may be biased, we suggest refined estimators to reduce the bias. To the best of our knowledge, a statistical approach has not been taken previously in the literature on applications of $L = \lambda W$ with measurements; e.g., see Denning and Buzen (1978), Little and Graves (2008), Little (2011), Lovejoy and Desmond (2011) and Mandelbaum (2011).

1.2.1. A Stationary Framework Two very different settings can arise: stationary and non-stationary. Preliminary data analysis should be done to determine if the data are from a stationary environment. In a stationary framework, we assume that Little’s law theory applies, so that L , λ and W are well defined, corresponding to both means of stationary probability distributions and limits of averages (assumed to exist), and related by $L = \lambda W$. We thus regard the underlying parameters L , λ and W as the *true* values that we want to estimate; we regard the averages $\bar{L}(t)$, $\bar{\lambda}(t)$ and $\bar{W}(t)$ based on measurements over a time interval $[0, t]$ as estimates of these parameters.

To learn how well we know L , λ and W when we compute the averages $\bar{L}(t)$, $\bar{\lambda}(t)$ and $\bar{W}(t)$, we suggest estimating confidence intervals. Given a single sample path from an interval that can be regarded as approximately stationary, we suggest applying the method of batch means to estimate confidence intervals, as is commonly done in simulation output analysis, and has been studied extensively; e.g., see Alexopoulos and Goldsman (2004), Asmussen and Glynn (2007), Tafazzoli et

al. (2011), Tafazzoli and Wilson (2011) and references therein. We present theory supporting its application in the present context.

In addition, we are concerned with the statistical problem of how to make inferences from limited data. We illustrate by focusing on estimating W given the finite averages $\bar{L}(t)$ and $\bar{\lambda}(t)$ when the waiting times are not directly observed. We pay special attention to the indirect estimator $\bar{W}_{L,\lambda}(t) \equiv \bar{L}(t)/\bar{\lambda}(t)$ suggested by Little's law. We show the special definition used to obtain equality for $\bar{L}(t) = \bar{\lambda}(t)\bar{W}(t)$ within each subinterval seriously distorts the batch-means estimators when the modified definition is used within each subinterval.

1.2.2. A Nonstationary Framework However, many applications with data involve nonstationary settings; e.g., service systems typically have arrival rates that vary significantly over each day. Estimation is more complicated without stationarity, because conventional Little's law theory no longer applies. Indeed, the parameters L , λ and W are typically no longer well defined. To specify what we are trying to estimate, we assume that there is an unspecified underlying stochastic queueing model, which may be highly nonstationary (for which the processes in §2.1 are well defined). As usual with Little's law, it is not necessary to define the underlying queueing model in detail. Then we regard the vector of time averages $(\bar{L}(t), \bar{\lambda}(t), \bar{W}(t))$ as a random vector with an associated vector of finite mean values $(E[\bar{L}(t)], E[\bar{\lambda}(t)], E[\bar{W}(t)])$. We propose that mean vector as the quantity to be estimated.

Since the method of batch means is no longer appropriate without stationarity, we suggest an approach corresponding to independent replications. That approach is appropriate for call centers when the data comes from multiple days that can be regarded as independent and identically distributed. In a nonstationary setting, the bias can be much more important, so we discuss ways to reduce it.

1.2.3. Validation by Simulation Since actual system data may be complicated and limited, we suggest applying simulation to study how the estimation procedures proposed here work for an idealized queueing model of the system. In doing so, we presume that we do not know enough

about the actual system to construct a model that we can directly apply to compute what we are trying to estimate, but that we know enough to be able to construct an idealized model to evaluate how the proposed estimation procedures perform. We illustrate this suggested simulation approach with our call center example in §3.2.

1.3. Organization

Here is how the rest of this paper is organized: In §2 we discuss the finite-time version of $L = \lambda W$, emphasizing the interval edge effects. In §3 we apply the statistical approach to a banking call center example and associated simulation models. In §4 we study ways to estimate confidence intervals. In §5 we study ways to estimate and reduce the bias in the estimator $\bar{W}_{L,\lambda}(t) \equiv \bar{L}(t)/\bar{\lambda}(t)$. In §6 we perform experiments combining the insights in §§4 and 5; we estimate confidence intervals for refined estimators designed to reduce bias. Finally, in §7 we draw conclusions. Additional material appears in the e-companion and a technical report (Kim and Whitt 2012) is available on the authors' web pages; the contents of both are described at the beginning of the e-companion.

2. Measurement over a Finite Time Interval: Definitions and Relations

In this section we review analogs of $L = \lambda W$ for a finite time interval, denoted by $[0, t]$. Consistent with most applications, we assume that the system was in operation in the past, prior to time 0, and that it will remain in operation after time t . We will use standard queueing terminology, referring to the items being counted as customers. We focus on the customers that are in the system at some time during the interval $[0, t]$. Let these customers be indexed in order of their arrival time, which could be prior to time 0 if the system is not initially empty (with some arbitrary method to break ties, if any).

2.1. The Performance Functions and Their Averages

For customer k , let A_k be the arrival time, D_k the departure time and $W_k \equiv D_k - A_k$ the waiting time (time in system), where $-\infty < A_k < D_k < \infty$, $[0, t] \cap [A_k, D_k] \neq \emptyset$ and \equiv denotes “equality by definition.” Let $R(0)$ count the customers that arrived before time 0 that remain in the system at

time 0; let $A(t)$ count the total number of new arrivals in the interval $[0, t]$; and let $L(t)$ be the number of customers in the system at time t . Thus, $A(t) = \max \{k \geq 0 : A_k \leq t\} - R(0)$, $t \geq 0$, and $L(0) = R(0) + A(0)$, where $A(0)$ is the number of new arrivals at time 0, if any. We will carefully distinguish between $R(0)$ and $L(0)$, but the common case is to have $A(0) = 0$ and $L(0) = R(0)$.

The respective averages over the time interval $[0, t]$ are

$$\bar{\lambda}(t) \equiv t^{-1}A(t), \quad \bar{L}(t) \equiv t^{-1} \int_0^t L(s) ds, \quad \bar{W}(t) \equiv (1/A(t)) \sum_{k=R(0)+1}^{R(0)+A(t)} W_k, \quad (1)$$

where $0/0 \equiv 0$ for $\bar{W}(t)$. The first two are time averages, while the last, $\bar{W}(t)$, is a customer average, but over all arrivals during the interval $[0, t]$.

We will focus on these averages over $[0, t]$ in (1), but we could equally well consider the averages associated with the first n arrivals. To do so, let T_n be the arrival epoch of the n^{th} new arrival, i.e., $T_n \equiv A_{n+R(0)}$, $n \geq 0$,

$$\bar{\lambda}_n \equiv n/T_n, \quad \bar{L}_n \equiv (1/T_n) \int_0^{T_n} L(s) ds, \quad \bar{W}_n \equiv n^{-1} \sum_{k=R(0)+1}^{R(0)+n} W_k. \quad (2)$$

As in (1), the first two averages in (2) are time averages, but over the time interval $[0, T_n]$, while the last, \bar{W}_n , is a customer average over the first n (new) arrivals. If there is only a single arrival at time T_n , then the averages in (2) can be expressed directly in terms of the averages in (1): $\bar{\lambda}_n = \bar{\lambda}(T_n)$, $\bar{L}_n = \bar{L}(T_n)$ and $\bar{W}_n = \bar{W}(T_n)$, so that conclusions for (1) yield analogs for (2).

Just as we can use the relation $L = \lambda W$ and knowledge of any two of the three quantities L , λ and W to compute the remaining one, so can we use any two of the three estimators in (1) to create a new alternative estimator, exploiting $L = \lambda W$:

$$\bar{L}_{W,\lambda}(t) \equiv \bar{\lambda}(t)\bar{W}(t), \quad \bar{\lambda}_{L,W}(t) \equiv \frac{\bar{L}(t)}{\bar{W}(t)} \quad \text{and} \quad \bar{W}_{L,\lambda}(t) \equiv \frac{\bar{L}(t)}{\bar{\lambda}(t)}. \quad (3)$$

For the typical application mentioned in §1 in which we observe $L(s)$, $0 \leq s \leq t$, we can directly construct the averages $\bar{L}(t)$ and $\bar{\lambda}(t)$, but we may not observe the individual waiting times. Hence, we may want to use $\bar{W}_{L,\lambda}(t)$ in (3) as a substitute for $\bar{W}(t)$ in (1).

2.2. How the Averages in (1) Are Related

Figures 1 and 2 below show how the three averages in (1) are related. These averages are related via $\bar{L}(t) = \bar{\lambda}(t)\bar{W}(t)$ if the system starts and ends empty, i.e., if $R(0) = L(t) = 0$, as we show in Theorem 1 below. However, more generally, these averages are not simply related. To illustrate, in Figures 1 and 2 a bar of height 1 is included for each of the customers in the system at some time during $[0, t]$ with the bar extending from the customer's arrival time to its departure time. (In this example the customers do not depart in the same order they arrived.) Thus the width of the bar is the customer's waiting time. For $0 \leq s \leq t$, the number of bars above any time s is $L(s)$.

To better communicate what is going on visually, we have ordered the customers in a special way. In Figures 1 and 2, the customers that arrive before time 0 but are still there at time 0 are placed first, starting at the bottom and proceeding upwards. These customers are ordered according to the arrival time, so the customers that arrived before time 0 appear at the bottom. One of these customers also departs after time t . The customers that arrived before time 0 and are still in the system at time 0 contribute to the regions A , B and C in Figure 2.

After the customers that arrived before time 0, we place the customers that arrive after time 0 and depart before time t , in order of arrival; they constitute region D in Figure 2. Finally, we place the customers that arrive after time 0 but depart after time t . These customers are ordered according to their arrival time as well; they constitute regions E and F in Figure 2. Three extra horizontal lines are included, along with the vertical lines at times 0 and t , to separate the regions. The arrival numbers are indicated along the vertical y axis. The condition $R(0) = L(t) = 0$ arises in Figure 2 as the special case in which all regions except region D are empty.

The averages can be expressed in terms of the two *cumulative processes*,

$$C_L(t) \equiv \int_0^t L(s) ds \quad \text{and} \quad C_W(t) \equiv \sum_{k=R(0)+1}^{R(0)+A(t)} W_k, \quad t \geq 0. \quad (4)$$

The difference between these two cumulative processes can be expressed in terms of the process $T_W^{(r)}(t)$, recording the *total residual waiting time* of all customers in the system at time t , i.e.,

$$T_W^{(r)}(t) \equiv \sum_{k=1}^{L(t)} W_k^{r,t}, \quad (5)$$

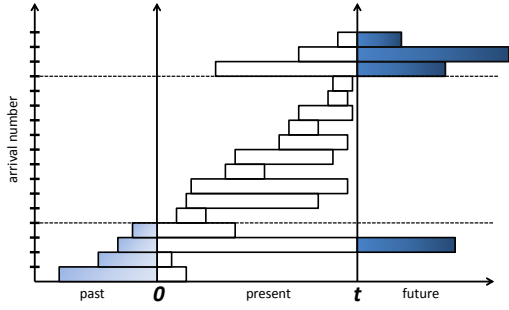


Figure 1 The total work in the system during the interval $[0, t]$ with edge effects: including arrivals before time 0 and departures after time t .

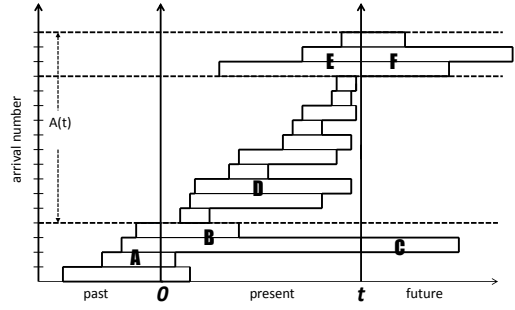


Figure 2 Six regions: waiting times (i) of customers that both arrive and depart inside $[0, t]$ (D), (ii) of arrivals before time 0 ($A \cup B \cup C$) and (iii) of departures after time t ($C \cup E \cup F$).

where $W_k^{r,t}$ is the remaining waiting time at time t for customer k in the system at time t (with index k assigned at time t among those remaining). The averages in (1) are the *time average* $\bar{L}(t) \equiv t^{-1}C_L(t)$ and the *customer average* $\bar{W}(t) \equiv C_W(t)/A(t)$. For a region A in Figure 2, let $|A|$ be the area of A . In general, the cumulative processes can be expressed in terms of the regions in Figure 2 as $C_L(t) = |B \cup D \cup E|$ and $C_W(t) = |D \cup E \cup F|$, while $T_W^{(r)}(0) = |B \cup C|$ and $T_W^{(r)}(t) = |C \cup F|$, so that

$$C_L(t) - C_W(t) = |B| - |F| = |B \cup C| - |F \cup C| = T_W^{(r)}(0) - T_W^{(r)}(t). \quad (6)$$

This relation for $C_L(t)$ is easy to see if we let ν be the total number of arrivals and departures in the interval $[0, t]$, τ_k be the k^{th} ordered time point among all the arrival times and departure times in $[0, t]$, with ties indexed arbitrarily and consistently, $\tau_0 \equiv 0$ and $\tau_{\nu+1} = t$. Then

$$C_L(t) \equiv \int_0^t L(s) ds = \sum_{j=1}^{\nu+1} \int_{\tau_{j-1}}^{\tau_j} L(s) ds = \sum_{j=1}^{\nu+1} L(\tau_{j-1})(\tau_j - \tau_{j-1}) = |B \cup D \cup E|,$$

where the last relation holds because $L(\tau_{j-1})$ is the number of single-customer unit-height bars above the interval $[\tau_{j-1}, \tau_j]$. Since $C_L(t) = C_W(t) = |D|$ if $R(0) = L(t) = 0$, we necessarily have the following well known result, appearing as Theorem I of Jewell (1967).

THEOREM 1. (*traditional finite-time Little's law*) If $R(0) = L(t) = 0$, then $\bar{L}(t) = \bar{\lambda}(t)\bar{W}(t)$.

Proof. Under the condition, $\bar{L}(t) \equiv \frac{C_L(t)}{t} = \frac{C_W(t)}{t} = \left(\frac{A(t)}{t}\right) \left(\frac{C_W(t)}{A(t)}\right) \equiv \bar{\lambda}(t)\bar{W}(t)$. ■

On the other hand, for the common case in which there are customers in the system during $[0, t]$ that arrived before time 0 and/or depart after time t , as in Figures 1 and 2, there is no simple relation between these cumulative processes and the associated averages, because of the interval edge effects. Nevertheless, the analysis above exposes the relationship that does hold. Variants of these relations are needed to establish sample-path limits in Little law theory, so the following result should not be considered new; e.g., see Theorem 1 of Glynn and Whitt (1986). A variant appears on p. 17.4 of Mandelbaum (2011), who credits it to his student Abir Koren and emphasizes its importance for looking at data.

THEOREM 2. (*extended finite-time Little's law*) The averages in (1) and (3) are related by

$$\begin{aligned}\Delta_L(t) &\equiv \bar{L}_{W,\lambda}(t) - \bar{L}(t) = \frac{|F| - |B|}{t} = \frac{T_W^{(r)}(t) - T_W^{(r)}(0)}{t}, \\ \Delta_W(t) &\equiv \bar{W}_{L,\lambda}(t) - \bar{W}(t) = \frac{|B| - |F|}{A(t)} = -\frac{\Delta_L(t)}{\bar{\lambda}(t)} = \frac{T_W^{(r)}(0) - T_W^{(r)}(t)}{A(t)}, \\ \Delta_\lambda(t) &\equiv \bar{\lambda}_{L,W}(t) - \bar{\lambda}(t) = \left(\frac{|B| - |F|}{|D| + |E| + |F|}\right) \bar{\lambda}(t) = -\frac{\Delta_L(t)}{\bar{W}(t)},\end{aligned}\tag{7}$$

where $|B|$ is the area of the region B in Figure 2 and $T_W^{(r)}(t)$ is defined in (5).

Since we focus on inferences about the average wait based on $\bar{L}(t)$ and $\bar{\lambda}(t)$ using $\bar{W}_{L,\lambda}(t)$, we focus on $\Delta_W(t)$ in (7). Given the customers need not depart in the order they arrive and we only observe $L(s)$, $0 \leq s \leq t$, the random variables $T_W^{(r)}(0)$ and $T_W^{(r)}(t)$ appearing in $\Delta_W(t)$ in (7) are not directly observable; we only have partial information about these random variables.

2.3. Alternative Definitions to Force Equality: The Inside View

Denning and Buzen (1978), Little (2011) and others have observed that we can preserve the relation $\bar{L}(t) = \bar{\lambda}(t)\bar{W}(t)$ in Theorem 1 without any conditions on $R(0)$ and $L(t)$ if we change the definitions. Equality can be achieved in general if we assume that our entire view of the system is *inside* the interval $[0, t]$. We see arrivals before time 0 but only as arrivals appearing at time 0, and we see the

portions of all waiting times only within the interval $[0, t]$. To achieve the inside view, let $A^{(i)}(t)$ count the number of new arrivals *plus* the number of customers initially in the system and let $W_k^{(i)}$ measure the waiting time *inside* the interval $[0, t]$; i.e., let

$$A^{(i)}(t) \equiv R(0) + A(t), \quad t \geq 0, \quad \text{and} \quad W_k^{(i)} \equiv (D_k \wedge t) - (A_k \vee 0), \quad k \geq 1, \quad (8)$$

where $a \wedge b \equiv \min\{a, b\}$ and $a \vee b \equiv \max\{a, b\}$. Now consider the associated averages

$$\bar{\lambda}^{(i)}(t) \equiv t^{-1} A^{(i)}(t) \quad \text{and} \quad \bar{W}^{(i)}(t) \equiv \frac{\sum_{k=1}^{A^{(i)}(t)} W_k^{(i)}}{A^{(i)}(t)}. \quad (9)$$

By an elementary modification of the proof of Theorem 1, we obtain the following “operational analysis” relation. (The equality relation corresponds to the operational Little’s law on p. 235 of Denning and Buzen (1978) and Theorem LL.2 of Little (2011).)

THEOREM 3. (*finite-time version of Little’s law with altered definitions*) *With the new definitions in (8) and (9), $\bar{\lambda}^{(i)}(t) \geq \bar{\lambda}(t)$, $\bar{W}^{(i)}(t) \leq \bar{W}(t)$ and $\bar{L}(t) = \bar{\lambda}^{(i)}(t)\bar{W}^{(i)}(t)$.*

Given that we only see inside the interval $[0, t]$, the reduced waiting times are *censored*. Indeed, there is no valid upper bound on $\bar{W}(t)$ based on the inside view. Arrivals before time 0 can have occurred arbitrarily far in the past prior to time 0, and customers present at time t can remain arbitrarily far into the future after time t . Any further properties of $\bar{W}(t)$ must depend on additional assumptions about what happens outside the interval $[0, t]$.

Even though the new definitions provide a good framework for checking the consistency of the data processing, and can be regarded as proper definitions, we advocate *not* using this modification because it causes problems in interpretation. We think it is usually better to account for the fact that an important part of the story takes place *outside* the interval $[0, t]$, even if we do not see it all. The alternative definitions in (8) also cause problems with the method of batch means used to construct confidence intervals; see §3.4.

3. A Banking Call Center Example

We illustrate the statistical approach by considering data from a telephone call center of an American bank from the data archive of Mandelbaum (2012). In 2001, this banking call center had sites in four states, which were integrated to form a single virtual call center. The virtual call center had 900 – 1200 agent positions on weekdays and 200 – 500 agent positions on weekends. The center processed about 300,000 calls per day during weekdays, with about 60,000 (20%) handled by agents, with the rest being served by integrated voice response (IVR) technology. As in many modern call centers, in this banking call center there were multiple agent types and multiple call types, with a form of skill-based routing (SBR) used to assign calls to agents.

Since we are only concerned with estimation related to the three parameters L , λ and W , we do not get involved with the full complexity of this system. For this paper, we use data for one class of customers, denoted by “Summit,” for 18 weekdays in May 2001; the data used and the analysis procedure are available from the authors’ web sites. Each working day covers a 17-hour period from 6 am to 11 pm, referred to as [6,23].

3.1. Sample Paths for a Typical Day

For some of the analyses, we will use a single day, Friday, May 25, 2001. Over this 17-hour period on that one day there were 5749 call arrivals (of this one type requesting an agent), of which 253 (4.4%) abandoned from queue before starting service. We do not include these abandonments in our analysis. Figures 3 and 4 show plots of the total number of arrivals into the queue (system), $A_q(s)$, and into service, $A_{ser}(s)$, together with the total number of departures from the queue (system), $D_q(s)$, and from service, $D_{ser}(s)$, all over the interval $[0, s]$, $0 \leq s \leq t$, first over the entire working day [6,23] and then over the hour [14,15]. These are based on the counts over 1-second subintervals. Note that the four curves in Figure 3 are too close to discern due to short waiting time (time in system, measured in minutes) relative to the time scale (hours). We see better when we zoom in, as in Figure 4.

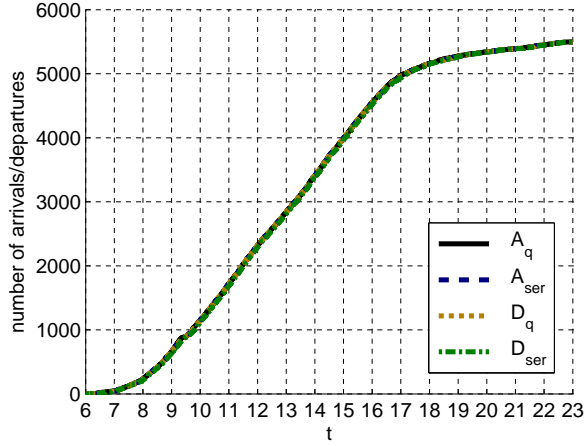


Figure 3 Arrivals and departures over the full day of May 25, 2001.

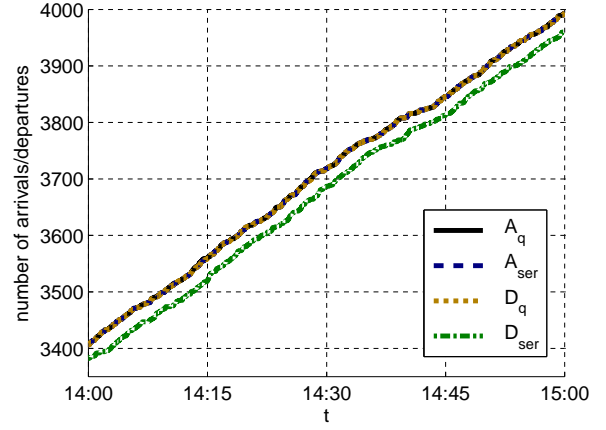


Figure 4 Arrivals and departures over the hour [14,15] within that day.

From the first plot in Figure 3, we see that the arrival rate is *not* stationary over the entire day (because the slope is not nearly constant), but it appears to be approximately stationary over the middle part of the day, e.g., in the six-hour interval [10,16]. When the arrival rate is nearly constant, so is the departure rate. The stationary-and-independent-increments property associated with a homogeneous Poisson process over [10,16] and the nonstationarity over [6,10] and [16,23] were confirmed by applying the turning points test, the difference-sign test and the rank test for randomness discussed on p. 312 of Brockwell and Davis (1991); the details appear in EC.2.

To confirm what we deduce from the arrival and departure rates, we also plot the number in system L_{sys} and the waiting times (times spent in the system), W_{sys} , and their hourly averages over the full day in Figures 5 and 6. Consistent with the plots in Figure 3, we see that the number in system looks approximately stationary in the 6-hour interval [10,16], but not over the full day [6,23]. In addition, Figure 6 shows that the hourly averages of the waiting times do not change much, especially in the interval [10,16]. During that 6-hour period [10,16], during which the system is approximately stationary, agents handled 3427 calls, of which only 28 (0.82%) abandoned. However, closer examination shows that the sample means \bar{L} are 28.3 and 32.6 over the hours [13,14] and [14,15], respectively, so that the differences can be shown to be statistically significant, but are minor compared to differences at the ends of the day. Since stationarity clearly does not

hold exactly, caution should be used in using the estimates.

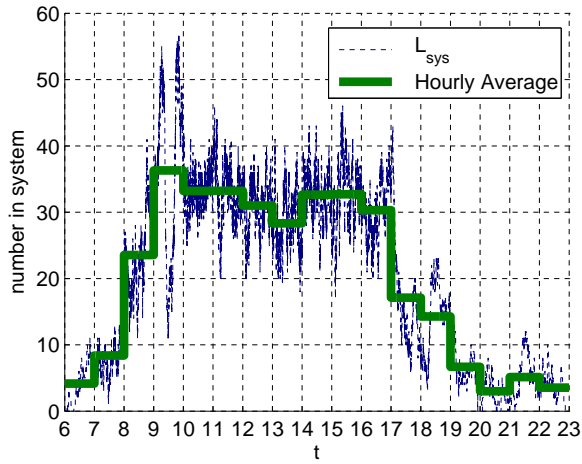


Figure 5 L_{sys} and its hourly averages over the full day.

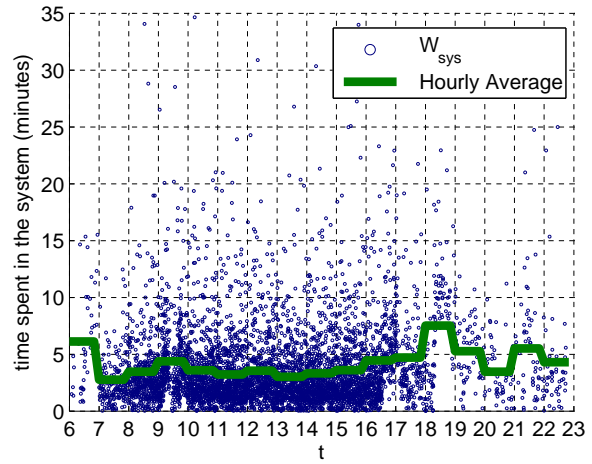


Figure 6 W_{sys} and its hourly averages over the full day.

To illustrate both the statistical approach to this example and the consequence of nonstationarity, we estimated L , λ and W both over the full day [6,23] and over the approximately stationary subinterval [10,16]. For both, we used the method of batch means, dividing the interval into $m = 20$ batches of equal length, producing batch lengths of 51 and 18 minutes, respectively. Over the full day, we have the estimates (measuring time in minutes)

$$\bar{L}_{full} = 20.2 \pm 6.1, \quad \bar{\lambda}_{full} = 5.39 \pm 1.84 \quad \text{and} \quad \bar{W}_{full} = 4.18 \pm 0.56; \quad (10)$$

over the interval [10,16], we have the estimates

$$\bar{L}_{stat} = 31.8 \pm 1.0, \quad \bar{\lambda}_{stat} = 9.44 \pm 0.31 \quad \text{and} \quad \bar{W}_{stat} = 3.39 \pm 0.15 \quad (11)$$

For each estimate in (10) and (11), we also include the halfwidth of the 95% confidence interval, estimated as described in the §4.3. We draw the following conclusions: (i) the confidence intervals tell us more than the averages alone, (ii) paying attention to stationarity is important, (iii) the halfwidths themselves reveal the nonstationarity, because we get far smaller halfwidths with the shorter subinterval [10,16], and (iv) since the mean waiting time is much less than the batch length, the number of batches is not grossly excessive (but that requires further examination).

3.2. Supporting Call Center Simulation Models

Many-server systems such as call centers are characterized by having many servers working independently in parallel. In such systems (if managed properly), the waiting times in queue tend to be short compared to the service times, and the service times tend to be approximately i.i.d. and independent of the arrival process. Thus, it is natural to use an idealized *infinite-server paradigm*, involving an infinite-server (IS) model with i.i.d. service times independent of the arrival process to approximately analyze statistical methods. Since the service times coincide exactly with the waiting times in the IS model, the waiting times are i.i.d. with constant mean $E[S]$, even though we are considering a nonstationary setting. That often holds approximately in service systems, as illustrated by our call center example.

For the call center, we have data on the arrival times and waiting times as well as the number in system $L(s)$, $0 \leq s \leq t$, but we do not have data on the staffing and the complex call routing. Thus, as suggested in §1.2.2, to evaluate the estimation procedures, we simulate the single-class single-service-pool $M_t/GI/\infty$ IS model and associated $M_t/GI/s_t$ models with time-varying staffing levels chosen to yield good performance, exploiting the square root staffing (SRS) formula $s(t) \equiv m(t) + \beta\sqrt{m(t)}$, where $m(t)$ is the *offered load*, the time-varying mean number of busy servers in the IS model, as in Jennings et al. (1996). As described in the Appendix EC.3.1, we fit the arrival rate function to a continuous piecewise-linear function, with one increasing piece over $[6, 10]$ starting at 0, a constant piece over $[10, 16]$ and two decreasing linear pieces over $[16, 18]$ and $[18, 23]$, the first steeper and the second ending at 0. We then simulated a nonhomogeneous Poisson arrival process with this arrival rate function. We assumed that all the service times were i.i.d. with a distribution obtained to match the observed waiting time distribution. A lognormal distribution with mean 3.38 and squared coefficient of variation $c_s^2 = 1.02$ was found to be a good fit, but an exponential distribution with that mean (and $c_s^2 = 1$) was also a good approximation, and so was used, because it is easier to analyze (see EC.3.2). The IS model was simulated with that fitted arrival rate function and service-time distribution. The offered load $m(t)$ was also computed by formulas (6)

and (7) of Jennings et al. (1996), drawing on Eick et al. (1993b); then the staffing function $s(t)$ was determined by the SRS formula using a range of Quality-of-Service (QoS) parameters β (see EC.3.3). We simulated 1000 independent replications of each of these models to study how the methods to estimate confidence intervals performed. In the next subsection we report results from simulation experiments showing that the finite-server models perform much like the IS model.

3.3. Confidence Intervals for the Call Center Data and Simulation

We applied the method of batch means to estimate confidence intervals for the parameters L , λ and W using the direct sample averages from (1) plus indirect estimate $\bar{W}_{L,\lambda}(t)$ from (3) for the time interval $[10, 16]$ over which the system is approximately stationary. (For both the call center data and the simulation model, we observe the waiting times, but we examine the alternative estimator $\bar{W}_{L,\lambda}(t)$ from (3) to see how it would perform if we could not observe the waiting times.)

We also consider the idealized $M_t/M/\infty$ and $M_t/M/s_t$ simulation models introduced in §3.2 and explained in detail in §EC.3. The estimation results are shown in Table 1. Additional results with more values of m appear in Tables 4-9 of Kim and Whitt (2012).

Table 1 Direct estimates of L , λ and W from (1) plus indirect estimate $\bar{W}_{L,\lambda}(t)$ from (3) with associated 95% confidence intervals for the approximately stationary time interval $[10, 16]$, constructed using batch means for $m = 5, 10$, and 20 batches for the call center data and idealized simulation models, including the $M_t/M/\infty$ and $M_t/M/s_t$ models with piece-wise linear arrival rate function fit to data, mean service time of 3.38 minutes and time-varying staffing based on the square-root-staffing formula using QoS parameter β taking values ranging from 1.0 to 2.5. Estimated confidence interval coverage is shown for the two waiting time estimates for the simulations based on 1000 replications.

case	m	$\bar{L}^{(m)}(t)$	$\bar{\lambda}^{(m)}(t)$	$\bar{W}^{(m)}(t)$	cov.	$\bar{W}_{L,\lambda}^{(m)}(t)$	cov.
$\beta = \infty$ ($M_t/M/\infty$)	5	31.5 ± 2.0	9.33 ± 0.42	3.38 ± 0.15	95.1%	3.38 ± 0.15	95.4%
	10	31.5 ± 1.6	9.33 ± 0.35	3.38 ± 0.13	95.0%	3.38 ± 0.13	95.7%
	20	31.5 ± 1.4	9.33 ± 0.33	3.38 ± 0.12	94.4%	3.38 ± 0.12	95.3%
$\beta = 2.5$ ($M_t/M/s_t$)	5	31.5 ± 2.0	9.33 ± 0.42	3.38 ± 0.15	95.3%	3.38 ± 0.15	95.9%
	10	31.5 ± 1.6	9.33 ± 0.35	3.38 ± 0.13	95.2%	3.38 ± 0.13	95.8%
	20	31.5 ± 1.4	9.33 ± 0.33	3.38 ± 0.12	95.0%	3.38 ± 0.12	95.3%
$\beta = 2.0$	5	31.5 ± 2.0	9.33 ± 0.42	3.38 ± 0.16	95.2%	3.38 ± 0.16	95.7%
	10	31.5 ± 1.6	9.33 ± 0.35	3.38 ± 0.13	95.3%	3.38 ± 0.13	95.6%
	20	31.5 ± 1.4	9.33 ± 0.33	3.38 ± 0.12	95.0%	3.38 ± 0.12	95.5%
$\beta = 1.5$	5	31.6 ± 2.2	9.33 ± 0.42	3.39 ± 0.17	95.8%	3.39 ± 0.17	95.9%
	10	31.6 ± 1.7	9.33 ± 0.35	3.39 ± 0.14	94.9%	3.39 ± 0.14	95.1%
	20	31.6 ± 1.5	9.33 ± 0.33	3.39 ± 0.13	94.0%	3.40 ± 0.13	94.9%
$\beta = 1.0$	5	32.1 ± 2.6	9.33 ± 0.42	3.44 ± 0.21	95.0%	3.44 ± 0.21	95.3%
	10	32.1 ± 2.1	9.33 ± 0.35	3.44 ± 0.17	93.2%	3.44 ± 0.17	93.5%
	20	32.1 ± 1.8	9.33 ± 0.33	3.44 ± 0.15	91.4%	3.44 ± 0.15	92.5%
data (call center)	5	31.9 ± 1.9	9.44 ± 0.49	3.38 ± 0.22		3.38 ± 0.19	
	10	31.9 ± 1.3	9.44 ± 0.36	3.39 ± 0.15		3.38 ± 0.16	
	20	31.9 ± 1.0	9.44 ± 0.30	3.39 ± 0.15		3.38 ± 0.11	

For large QoS parameter β , e.g., $\beta \geq 2.0$, the performance in the finite-server model is essentially the same as in the associated IS model, as can be seen from Table 1. However, as β decreases, more customers have to wait before starting service. Thus, the estimated mean waiting time increases from 3.38 in the IS model to 3.39 and 3.44, respectively, for $\beta = 1.5$ and 1.0, respectively. Similarly, the estimated mean number in system increases from 31.5 to 31.6 and 32.1 for these same cases. Of special interest is the confidence interval coverage in the simulations based on 1000 replications. Table 1 shows it is excellent for all values of m , being very close to the target 95.0%, for all $\beta \geq 1.5$. However, we see a drop in coverage for $\beta = 1$. Thus, to be conservative, we advocate using for the call center model the largest estimated CI, which usually should be associated with the smallest number of batches $m = 5$ for the call center data. Overall, Table 1 shows that the indirect estimator $\bar{W}_{L,\lambda}^{(m)}(t)$ behaves very much the same as the direct estimator $\bar{W}(t)$. Indeed, that is consistent with the theory and other experiments in this paper.

To illustrate what happens with a shorter sample path segment, we consider the interval [14, 15].

Table 2 shows the corresponding estimates for the IS model and the call center. Additional results with more values of m appear in Tables 10 and 11 of Kim and Whitt (2012). In this case, $m = 5$, 10, and 20 corresponds to 5 batches of 12 minutes, 10 batches of 6 minutes and 20 batches of 3 minutes, respectively. The CI coverage is again excellent for the IS model for all cases. However, since the mean waiting time is about 3.4 minutes, we regard only $m = 5$ appropriate for the interval $[14, 15]$ (We also note that the difference between $\bar{W}_{L,\lambda}^{(m)}(t)$ and $\bar{W}(t)$ in Table 2 becomes greater as m increases. See Section 3.4.1 for more discussion on this).

Table 2 Direct estimates of L , λ and W from (1) plus indirect estimate $\bar{W}_{L,\lambda}(t)$ from (3) with associated 95% confidence intervals for the approximately stationary time interval $[14, 15]$ constructed using batch means for $m = 5, 10$, and 20 batches for the call center data and simulation of the idealized $M_t/M/\infty$ models, with piecewise linear arrival rate function fit to data, mean service time of 3.38 minutes and time-varying staffing based on the square-root-staffing formula using QoS parameter β . Estimated confidence interval coverage based on 1000 replications is shown for the two waiting time estimates for the simulations.

case	m	$\bar{L}(t)$	$\bar{\lambda}(t)$	$\bar{W}(t)$	cov.	$\bar{W}_{L,\lambda}^{(m)}(t)$	cov.
$\beta = \infty$ ($M_t/M/\infty$)	5	31.4 ± 4.0	9.32 ± 1.04	3.37 ± 0.37	95.6%	3.38 ± 0.37	94.8%
	10	31.4 ± 2.9	9.32 ± 0.87	3.37 ± 0.32	95.8%	3.40 ± 0.32	95.4%
	20	31.4 ± 2.1	9.32 ± 0.82	3.37 ± 0.30	95.9%	3.46 ± 0.32	94.3%
data (call center)	5	32.6 ± 1.9	9.82 ± 0.82	3.33 ± 0.21		3.33 ± 0.10	
	10	32.6 ± 1.6	9.82 ± 0.79	3.33 ± 0.21		3.34 ± 0.16	
	20	32.6 ± 1.3	9.82 ± 0.81	3.32 ± 0.23		3.43 ± 0.31	

3.4. Edge Effects and the Method of Batch Means

The issue of interval edge effects discussed in §2 becomes more serious with the method of batch means. For a fixed sample path segment of length t and m batches, there are m intervals, each with edge effects, and each interval is of length t/m instead of t .

3.4.1. The Error Due to the Interval Edge Effects Formula (7) shows that the difference between $\bar{W}_{L,\lambda}(t)$ and $\bar{W}(t)$ should be inversely proportional to t in a stationary setting, because the distribution of $T_W^{(r)}(t)$ is independent of t , whereas $\bar{\lambda}(t) \equiv t^{-1}A(t) \rightarrow \lambda$ as $t \rightarrow \infty$. We should expect serious bias if t is less than or equal to W , the average time spent in the system, but very little bias if t is much greater. Since $W \approx 3.4$ minutes for the call-center example from §3, we expect

serious bias if $t = 3$ minutes, some bias if $t = 30$ minutes and almost no bias if $t = 300$ minutes. Those expectations are confirmed by the averages shown in Table 3. In each case, the averages over subintervals correspond to batch means. (See Tables 12-14 of Kim and Whitt (2012) for more details.)

Table 3 Comparison of the direct and indirect estimators $\bar{W}(t)$ and $\bar{W}_{L,\lambda}(t)$ for three values of t : 3, 30 and 300 minutes for the data from §3. Averages are given for the 20 subintervals of $[14, 15]$ for $t = 3$ minutes, for the 20 subintervals of $[8, 18]$ for $t = 30$ minutes and for the 4 overlapping 5-hour subintervals of $[9, 17]$, from $[9, 14]$ to $[12, 17]$, for $t = 300$ minutes.

t	$\bar{W}(t)$	$\bar{W}_{L,\lambda}(t)$	$ \Delta_W(t) $	$\Delta_W^{rel}(t)$	$ U $	$ A $	$ B $	$ C $	$ D $	$ E $	$ F $	$ F - B $
3	3.32	3.43	0.713	20.3%	311	34.8%	19.5%	14.3%	3.1%	8.9%	19.4%	6.3%
30	3.80	3.80	0.241	5.6%	1101	9.7%	9.2%	0.0%	62.8%	9.4%	8.8%	4.5%
300	3.44	3.44	0.016	0.5%	9754	1.4%	1.2%	0.0%	94.9%	1.2%	1.3%	0.4%

In Table 3 we see that the relative error $\Delta_W^{rel}(t) \equiv \Delta_W(t)/\bar{W}_{L,\lambda}(t)$ takes the values 20.3%, 5.6% and 0.5%, respectively for $t = 3, 30$ and 300 minutes. For the regions in Figure 2, for $t \geq 30$ minutes, we see that $|C| = 0$, the areas of regions B , C , E and F are approximately independent of t , while the area of D is proportional to t . Table 3 shows the area of the union of all six regions, $U \equiv A \cup B \cup C \cup D \cup E \cup F$, and the percentages of that total area made up by each of the six regions, as well as $|F| - |B|$. The simple case occurs when region D dominates the six regions. The percentage of the total area provided by D is 94.9% for $t = 300$ minutes, 62.8% for $t = 30$ minutes and 3.1% for $t = 3$ minutes.

3.4.2. Additional Error from the Altered Definitions The altered definitions in §2.3 become more unattractive with batch means, because the shorter intervals distort the meaning even more. The average truncated waiting times $\bar{W}_c(t)$ in (9) tend to be even less than the true average waiting times W , while the average augmented arrivals $\bar{\lambda}_i(t)$ in (9) tend to be even more than the true average arrival rate λ . The altered definitions lead to double counting for arrivals. Customers that are in the system during more than one interval are counted as arrivals in all these intervals.

To illustrate, we consider the call center data over the interval $[10, 16]$. Without using batches, we have $\bar{\lambda}(t) = 9.44$ arrivals per minute and $\bar{W}(t) = 3.38$ minute, while the estimators using the altered definitions in (9) are $\bar{\lambda}_i(t) = 9.55$ and $\bar{W}_c(t) = 3.33$. With m batches, $1 \leq m \leq 20$, the estimator $\bar{\lambda}(t)$ is unchanged and the estimator $\bar{W}(t)$ differs by only 0.001 from the original value of 3.38 for $m = 1$. In contrast, $\bar{\lambda}_i(t)$ assumes the values 9.55, 9.88, 10.33 and 11.16 for $m = 1, 5, 10$ and 20, respectively. Similarly, $\bar{W}_c(t)$ assumes the values 3.33, 3.22, 3.09 and 2.86 for $m = 1, 5, 10$ and 20, respectively. For $m = 20$, the errors in $\bar{\lambda}_i(t)$ and $\bar{W}_c(t)$ are 18% and 15%, respectively. When confidence intervals are formed based on batch means (for non-negligible m), the systematic errors caused by the altered definition far exceed the halfwidth of the confidence interval. Hence we recommend not using the modified definitions in (9).

4. Confidence Intervals: Theory and Methodology

We now consider how to apply the estimator $\bar{W}_{L,\lambda}(t)$ in (3) to estimate a confidence interval (CI) for W in a stationary setting and for $E[\bar{W}(t)]$ in a nonstationary setting, without observing the waiting times. We will be using statistical methods commonly used in simulation experiments. However, unlike simulation, we anticipate that system data is likely to be limited, so we may not be able to achieve high precision. Nevertheless, we want to have some idea how well we know the estimated values. With that in mind, we suggest applying standard statistical methods. In order to evaluate how well these statistical procedures should perform, e.g., to verify that CI coverage should be approximately as specified, we advocate studying associated idealized simulation models of the system more closely as suggested in §1.2.2, and as illustrated in §3.2.

For the common case in which we have only a single sample path segment, we advocate applying the method of batch means, as specified in §4.3. That method depends on the batch means being approximately i.i.d. and normally distributed. We point out that there is a risk that these assumptions may not be justified, so that estimated CI's should be used with caution. We suggest using multiple i.i.d. replications of the supporting simulation model to confirm these properties and evaluate the confidence interval coverage. If these standard methods do not perform well for

the supporting simulation models, then we can consider more sophisticated estimation methods, as in Alexopoulos et al. (2007), Tafazzoli et al. (2011), Tafazzoli and Wilson (2011) and references therein.

4.1. A Ratio Estimator

In both stationary and nonstationary settings, a CI (interval estimate) for $E[\bar{W}(t)]$ without observing the waiting times can be obtained using $\bar{W}_{L,\lambda}(t)$ if we can apply the following theorem, implementing the delta method; see §III.3 and Proposition §IV.4.1 of Asmussen and Glynn (2007) for related results.

THEOREM 4. (*asymptotics for the ratio of low-variability positive normal random variables*) *If there is a sequence of systems indexed by n such that*

$$\sqrt{n}(\bar{L}^{(n)}(t) - L, \bar{\lambda}^{(n)}(t) - \lambda) \Rightarrow N(0, \Sigma) \quad \text{in } \mathbb{R}^2 \quad \text{as } n \rightarrow \infty, \quad (12)$$

where L and λ are positive real numbers and $N(0, \Sigma)$ is a mean-zero bivariate Gaussian random vector with variance vector $(\sigma_L^2, \sigma_\lambda^2)$ and covariance $\sigma_{L,\lambda}^2$, and $\bar{W}^{(n)}(t)$ satisfies

$$\bar{W}^{(n)}(t)/\bar{W}_{L,\lambda}^{(n)}(t) \Rightarrow 1 \quad \text{as } n \rightarrow \infty, \quad (13)$$

for $\bar{W}_{L,\lambda}^{(n)}(t) \equiv \bar{L}^{(n)}(t)/\bar{\lambda}^{(n)}(t)$, then

$$\sqrt{n}(\bar{W}^{(n)}(t) - (L/\lambda)) \Rightarrow N(0, \sigma_W^2) \quad \text{in } \mathbb{R} \quad \text{as } n \rightarrow \infty \quad (14)$$

for

$$\sigma_W^2 = \frac{1}{\lambda^2} \left(\sigma_L^2 - \frac{2L\sigma_{L,\lambda}^2}{\lambda} + \frac{L^2\sigma_\lambda^2}{\lambda^2} \right). \quad (15)$$

Proof. Apply a Taylor expansion with the function $f(x, y) \equiv x/y$, having first partial derivatives $f_x = 1/y$ and $f_y = -x/y^2$, to get

$$\frac{\bar{L}^{(n)}(t)}{\bar{\lambda}^{(n)}(t)} = \frac{L}{\lambda} + \frac{\bar{L}^{(n)}(t) - L}{\lambda} - \frac{L(\bar{\lambda}^{(n)}(t) - \lambda)}{\lambda^2} + o(\max\{|\bar{L}^{(n)}(t) - L|, |\bar{\lambda}^{(n)}(t) - \lambda|\}), \quad (16)$$

so that

$$\sqrt{n} \left(\bar{W}_{L,\lambda}^{(n)}(t) - (L/\lambda) \right) = \frac{\sqrt{n}(\bar{L}^{(n)}(t) - L)}{\lambda} - \frac{\sqrt{n}L(\bar{\lambda}^{(n)}(t) - \lambda)}{\lambda^2} + o(1) \quad \text{as } n \rightarrow \infty, \quad (17)$$

from which (14) follows, given (12) and (13). ■

We can apply the theorem if our system can be regarded as system n for n sufficiently large that we can replace the limits with approximate equality. The approximate confidence interval estimate for $E[\bar{W}^{(n)}(t)]$ would then be $[\bar{W}_{L,\lambda}^{(n)}(t) - 1.96\sigma_W/\sqrt{n}, \bar{W}_{L,\lambda}^{(n)}(t) + 1.96\sigma_W/\sqrt{n}]$, where σ_W is the square root of the variance σ_W^2 in (15). Since the variance σ_W^2 in (15) is typically unknown, we must estimate it. That can be done by inserting estimates for all the components of (15). Assuming that the estimates converge as $n \rightarrow \infty$, we still have asymptotic normality with the estimated values of the variance σ_W^2 .

The sequence of systems indexed by n satisfying condition (12) in Theorem 4 can arise in two natural ways: First, condition (12) is typically satisfied if the averages are collected from a single observation over successively longer time intervals in a stationary environment, i.e., if t is allowed to grow with n , with $t_n \rightarrow \infty$. Then, of course, $E[\bar{W}^{(n)}(t)] \rightarrow W$ as $n \rightarrow \infty$, and we are simply estimating W . Second, whether or not there is a stationary environment, condition (12) is satisfied if the averages indexed by n correspond to averages taken over n multiple independent samples for a fixed interval $[0, t]$. The second case is important for the common case of service systems with strongly time-varying arrival rates over each day, provided that multiple days can be regarded as i.i.d. samples.

Condition (13) in Theorem 4 is of course also satisfied if the averages are collected from a single observation over successively longer time intervals in a stationary environment. However, condition (13) may well *not* be satisfied, even approximately, if the averages indexed by n correspond to averages taken over n multiple independent samples for a fixed interval $[0, t]$, because the bias may be significant, and it does not go away with increasing n ; see §5.

4.2. The Supporting Central Limit Theorem in a Stationary Setting

With one sample path segment, we suggest applying the method of batch means. A partial basis for that is the central limit theorem (CLT) version of Little's law in Glynn and Whitt (1986) and Whitt (2012). To apply it, we assume that the system is approximately stationary over the designated subinterval $[0, t]$. Hence we regard the finite averages in (1) as estimators of the unknown parameters L , λ and W . The CLT states that, under very general regularity conditions,

$$(\hat{L}(t), \hat{\lambda}(t), \hat{W}(t), \hat{L}_{W,\lambda}(t), \hat{\lambda}_{L,W}(t), \hat{W}_{L,\lambda}(t)) \Rightarrow (X_L, X_\lambda, X_W, X_L, X_\lambda, X_W) \quad \text{in } \mathbb{R}^6 \quad (18)$$

as $t \rightarrow \infty$, where

$$\begin{aligned} (\hat{L}(t), \hat{\lambda}(t), \hat{W}(t)) &\equiv \sqrt{t} (\bar{L}(t) - L, \bar{\lambda}(t) - \lambda, \bar{W}(t) - W), \\ (\hat{L}_{W,\lambda}(t), \hat{\lambda}_{L,W}(t), \hat{W}_{L,\lambda}(t)) &\equiv \sqrt{t} (\bar{L}_{W,\lambda}(t) - L, \bar{\lambda}_{L,W}(t) - \lambda, \bar{W}_{L,\lambda}(t) - W), \end{aligned} \quad (19)$$

with the averages given in (1) and (3), and the limiting random vector (X_L, X_λ, X_W) is an essentially two-dimensional mean-zero multivariate Gaussian random vector with $X_W = \lambda^{-1}(X_L - W X_\lambda)$, so that the variance and covariance terms are related by

$$\begin{aligned} \sigma_W^2 &\equiv \text{Var}(X_W) = E[X_W^2] = \lambda^{-2}(\sigma_L^2 - 2W\sigma_{\lambda,L}^2 + W^2\sigma_\lambda^2), \\ \sigma_{L,W}^2 &\equiv \text{Cov}(X_L, X_W) = E[X_L X_W] = \lambda^{-1}(\sigma_L^2 - W\sigma_{\lambda,L}^2), \\ \sigma_{W,\lambda}^2 &\equiv \text{Cov}(X_W, X_\lambda) = E[X_W X_\lambda] = \lambda^{-1}(\sigma_{\lambda,L}^2 - W\sigma_\lambda^2). \end{aligned} \quad (20)$$

Note that σ_W^2 in (20) agrees with (15).

Under general regularity conditions (essentially, if $t^{1/2}T_W^{(r)}(t) \Rightarrow 0$ for $T_W^{(r)}(t)$ in (5)), a functional central limit theorem (FCLT) generalization of the joint CLT in (18) is valid if a FCLT is valid in \mathbb{R}^2 for any two of the first three components. For example, it suffices start with the (FCLT generalization of) the bivariate CLT

$$\sqrt{t}(\bar{L}(t) - L, \bar{\lambda}(t) - \lambda) \Rightarrow (X_L, X_\lambda) \quad \text{in } \mathbb{R}^2 \quad \text{as } t \rightarrow \infty, \quad (21)$$

where the limit (X_L, X_λ) is a bivariate mean-zero Gaussian random vector with variances σ_λ^2 , σ_L^2 and covariance $\sigma_{\lambda,L}^2$. Natural sufficient conditions are based on regenerative structure for the stochastic process $\{L(t) : t \geq 0\}$, as in §VI.3 of Asmussen (2003) and Glynn and Whitt (1987). We directly assume that the limit in (18) is valid, and discuss how to apply it. Note that condition (21) coincides with condition (12) in Theorem 4, but now the conclusion directly gives a CLT for $\bar{W}(t)$ as well as for $\bar{W}_{L,\lambda}(t)$.

The form of the limit in (18) implies that the alternative estimators $\bar{L}_{W,\lambda}(t)$, $\bar{\lambda}_{L,W}(t)$ and $\bar{W}_{L,\lambda}(t)$ in (3) not only converge to the same limits L , λ and W just as the natural estimators $\bar{L}(t)$, $\bar{\lambda}(t)$ and $\bar{W}(t)$ in (1) do, but also the corresponding CLT-scaled random variables are asymptotically equivalent as well, i.e., $\|(\hat{L}(t), \hat{\lambda}(t), \hat{W}(t)) - (\hat{L}_{W,\lambda}(t), \hat{\lambda}_{L,W}(t), \hat{W}_{L,\lambda}(t))\| \Rightarrow 0$ as $t \rightarrow \infty$, where $\|\cdot\|$ is the Euclidean norm on \mathbb{R}^3 .

In summary, the CLT version of $L = \lambda W$ implies that the asymptotic efficiency (halfwidth of confidence intervals for large sample sizes) is the same for the alternative estimators in (3) as it is for the natural estimators in (1) (in a stationary setting). However, if one of the parameters happens to be known in advance, one estimator can be more efficient than the other; see Glynn and Whitt (1989). For example, with simulation, the arrival rate is typically known in advance.

4.3. Estimating Confidence Intervals by the Method of Batch Means

Assuming that the conditions for the CLT in the previous section are satisfied, given the sample path segments $\{(A(s), L(s)) : 0 \leq s \leq t\}$ and $\{W_k : R(0) + 1 \leq k \leq R(0) + A(t)\}$ over the time interval $[0, t]$ (or only two of these three segments), we can use m batches based on measurements over the m subintervals $[(k-1)t/m, kt/m]$, $1 \leq k \leq m$. To define the batch averages, let $R_k \equiv R(kt/m)$, the number of customers remaining in the system at time kt/m from among those that arrived previously. Let $\bar{A}_k(t, m)$, $\bar{L}_k(t, m)$ and $\bar{W}_k(t, m)$ denote the averages over the interval $[(k-1)t/m, kt/m]$, i.e.,

$$\bar{A}_k(t, m) \equiv (m/t)A_k(t, m), \quad \bar{L}_k(t, m) \equiv (m/t)L_k(t, m), \quad \bar{W}_k(t, m) \equiv (1/A_k(t, m))W_k(t, m), \quad (22)$$

$$A_k(t, m) \equiv A(kt/m) - A((k-1)t/m), \quad L_k(t, m) \equiv \int_{(k-1)t/m}^{kt/m} L(s) ds, \quad W_k(t, m) \equiv \sum_{j=R_{k-1}+1}^{R_k-1+A_k(t, m)} W_j.$$

The FCLT version of the CLT in the previous section implies that, as $t \rightarrow \infty$, the vector of scaled batch means $\sqrt{t/m}(\bar{A}_k(t, m) - \lambda, \bar{L}_k(t, m) - L, \bar{W}_k(t, m) - W)$, $1 \leq k \leq m$, are asymptotically m i.i.d. mean-zero Gaussian random vectors with variances σ_λ^2 , σ_L^2 and σ_W^2 , and covariances $\sigma_{L,\lambda}^2$, $\sigma_{\lambda,W}^2$ and $\sigma_{L,W}^2$. By Theorem 4, as $t \rightarrow \infty$, the associated scaled vector $\sqrt{t/m}(\bar{W}_{L,\lambda,k}(t, m) - W)$, $1 \leq k \leq m$, are asymptotically m i.i.d. mean-zero random variables with variance σ_W^2 in (15). Hence, as $t \rightarrow \infty$, also

$$\frac{\sum_{k=1}^m (\bar{W}_{L,\lambda,k}(t, m) - \bar{W}_{L,\lambda}^{(m)}(t))}{\sqrt{S_{(m)}^2(t)/m}} \Rightarrow t_{m-1}, \quad (23)$$

where t_{m-1} is a random variable with the Student t distribution with $m-1$ degrees of freedom,

$$\bar{W}_{L,\lambda}^{(m)}(t) \equiv \frac{1}{m} \sum_{k=1}^m \bar{W}_{L,\lambda,k}(t, m) \quad \text{and} \quad S_{(m)}^2(t) \equiv \frac{1}{m-1} \sum_{k=1}^m (\bar{W}_{L,\lambda,k}(t, m) - \bar{W}_{L,\lambda}^{(m)}(t))^2. \quad (24)$$

Thus, $[\bar{W}_{L,\lambda}^{(m)}(t) - t_{0.025, m-1} S_{(m)}(t)/\sqrt{m}, \bar{W}_{L,\lambda}^{(m)}(t) + t_{0.025, m-1} S_{(m)}(t)/\sqrt{m}]$ is an approximate 95% confidence interval for W based on the t distribution and the average $\bar{W}_{L,\lambda}^{(m)}(t)$ of batch means. Of course the same procedure applies to other averages of batch means as well.

It remains to choose the number of batches, m . Since we obtain larger batch sizes, and thus more nearly approximate the asymptotic condition $t \rightarrow \infty$, if we make m small, we advocate keeping it relatively small, e.g., $m = 5$. Nevertheless, in our examples we consider a range of m values.

5. Estimating and Reducing the Bias

We now discuss ways to estimate and reduce the bias in the estimator $\bar{W}_{L,\lambda}(t)$ in (3) as an estimator for $E[\bar{W}(t)]$ for $\bar{W}(t)$ in (1). In doing so, we are primarily concerned with nonstationary settings. In stationary settings, $\bar{W}(t)$ in (1) is typically a biased estimator of W , while $\bar{W}_{L,\lambda}(t)$ is typically a biased estimator of both W and $E[\bar{W}(t)]$, but these biases are less likely to be serious, e.g., see §5.4.

An important conclusion from our analysis is that the bias depends on the underlying model. We demonstrate by considering two idealized paradigms: the infinite-server and single-server

paradigms. We emphasize the infinite-server paradigm, which often is appropriate for call centers. In §5.4, we show that the bias in $\bar{W}(t)$ for estimating W tends to be negligible in the infinite-server paradigm.

5.1. Bias in $\bar{W}_{L,\lambda}(t)$ as an Estimator of the Expected Average Wait $E[\bar{W}(t)]$

Since the bias in $\bar{W}_{L,\lambda}(t)$ as an estimator for $E[\bar{W}(t)]$ is $E[\Delta_W(t)]$ for $\Delta_W(t) \equiv \bar{W}_{L,\lambda}(t) - \bar{W}(t)$ in (7), we can apply Theorem 2 to obtain an exact expression for the bias $E[\Delta_W(t)]$. We also give the conditional bias $E[\Delta_W(t)|\mathcal{O}_t]$ given the observed data over the interval $[0, t]$, which we assume is $\mathcal{O}(t) \equiv (t, \bar{L}(t), \bar{\lambda}(t), R(0), L(t))$, from which we can also deduce $A(t)$. We use the conditional bias to create a refined estimator given the observed data.

COROLLARY 1. (*exact bias and conditional bias*) *The bias in $\bar{W}_{L,\lambda}(t)$ in (3) as an estimator for $E[\bar{W}(t)]$ for $\bar{W}(t)$ in (1) is $E[\Delta_W(t)] = E[E[\Delta_W(t)|\mathcal{O}(t)]]$, where $\Delta_W(t)$ is given in (7), the vector of observed data is $\mathcal{O}(t) \equiv (t, \bar{L}(t), \bar{\lambda}(t), R(0), L(t))$ and the conditional bias is*

$$E[\Delta_W(t)|\mathcal{O}_t] = \frac{\sum_{k=1}^{R(0)} E[W_k^{r,0}|\mathcal{O}_t] - \sum_{k=1}^{L(t)} E[W_k^{r,t}|\mathcal{O}_t]}{A(t)}. \quad (25)$$

Proof. Apply Theorem 2 using (5). ■

5.2. Two Approximations

The bias in Corollary 1 is not easy to analyze. Given that $(R(0), L(t), A(t))$ is observed, it remains to estimate the conditional residual waiting times $E[W_k^{r,0}|\mathcal{O}_t]$, $1 \leq k \leq R(0)$, and $E[W_k^{r,t}|\mathcal{O}_t]$, $1 \leq k \leq L(t)$. The conditional expectations $E[W_k^{r,0}|\mathcal{O}_t]$ are complicated, because we are conditioning on events in the future after the observation time 0. Thus, we develop two approximations and then show that they apply to the infinite-server paradigm.

5.2.1. Simplification from the Bias Approximation Assumption As t increases, we expect the “initial edge effect” $\{R(0), W_k^{r,0}; 1 \leq k \leq R(0)\}$ to be approximately independent of the “terminal edge effect” $\{L(t), W_k^{r,t}; 1 \leq k \leq L(t)\}$ and the total number of arrivals $A(t)$. With that in mind, we use the following approximation, which primarily means that we are assuming that t is sufficiently large.

Bias Approximation Assumption (BAA). For $\mathcal{O}(t) \equiv (t, \bar{L}(t), \bar{\lambda}(t), R(0), L(t))$, $t \geq 0$,

$$E[W_k^{r,0}|\mathcal{O}_t] \approx E[W_k^{r,0}|R(0)], \quad 0 \leq k \leq R(0), \quad \text{and} \quad E[W_k^{r,t}|\mathcal{O}_t] \approx E[W_k^{r,t}|L(t)], \quad 0 \leq k \leq L(t).$$

Invoking the BAA, we obtain the following approximation directly from Corollary 1:

$$E[\Delta_W(t)|\mathcal{O}_t] \approx \frac{\sum_{k=1}^{R(0)} E[W_k^{r,0}|R(0)] - \sum_{k=1}^{L(t)} E[W_k^{r,t}|L(t)]}{A(t)}. \quad (26)$$

We think that BAA is reasonable if t is sufficiently large. That is easy to see for stationary models, because then as $t \rightarrow \infty$ (i) $\bar{L}(t) \rightarrow L$ and $\bar{\lambda}(t) \rightarrow \lambda$ and (ii) under regularity conditions (e.g., regenerative structure), $\{R(0), W_k^{r,0}; 1 \leq k \leq R(0)\}$ will be asymptotically independent of $\{L(t), W_k^{r,t}; 1 \leq k \leq L(t)\}$.

5.2.2. Using $\bar{W}_{L,\lambda}(t)$ to Estimate the Residual Waiting Times We can obtain an applicable estimate of the conditional bias $E[\Delta_W(t)|\mathcal{O}_t]$ in (25) if we estimate all the remaining conditional waiting times by the observed $\bar{W}_{L,\lambda}(t)$. In doing so, we are ignoring the inspection paradox (since these are remainders of waiting times in progress), the model structure and the available information $\mathcal{O}(t)$. This step is likely to be justified approximately if the distribution of the waiting times is nearly exponential.

That step yields the approximation

$$E[\Delta_W(t)|\mathcal{O}_t] \approx \frac{(R(0) - L(t))\bar{W}_{L,\lambda}(t)}{A(t)} \quad \text{for} \quad \mathcal{O}(t) \equiv (R(0), L(t), \bar{L}(t), \bar{\lambda}(t)). \quad (27)$$

We can apply approximation (27) to obtain the new candidate *refined estimator* of $E[\bar{W}(t)]$, exploiting the observed vector $(R(0), L(t), A(t))$:

$$\bar{W}_{L,\lambda,r}(t) \equiv \bar{W}_{L,\lambda}(t) - E[\Delta_W(t)|\mathcal{O}_t] \approx \bar{W}_{L,\lambda}(t) \left(1 - \frac{R(0) - L(t)}{A(t)}\right). \quad (28)$$

(The refined estimator $\bar{W}_{L,\lambda,r}(t)$ in (28) is a candidate refinement of the indirect estimator $\bar{W}_{L,\lambda}(t)$

(3).) The associated approximate relative conditional bias is thus

$$E[\Delta_W^{rel}(t)|\mathcal{O}(t)] \equiv \frac{E[\Delta_W(t)|\mathcal{O}_t]}{E[\bar{W}(t)]} \approx \frac{E[\Delta_W(t)|\mathcal{O}_t]}{\bar{W}_{L,\lambda}(t)} \approx \frac{R(0) - L(t)}{A(t)}. \quad (29)$$

In the next section we show that the analysis in (27)-(29) can be supported theoretically in the infinite-server paradigm when the waiting times are exponential, so we propose the refined estimator in (28) as a candidate estimator for many-server systems. However, the crude analysis above is not justified universally; e.g., it is not good for the single-server models, as we show in §5.5.

5.3. The Infinite-Server Paradigm

If, in addition to BAA, we consider the $G_t/M/\infty$ IS model with exponential service times having mean $E[S]$, then (26) becomes

$$E[\Delta_W(t)|\mathcal{O}_t] \approx (R(0) - L(t))E[S]/A(t). \quad (30)$$

Since the waiting times coincide with the service times in the IS model, it is natural to use the observed $\bar{W}_{L,\lambda}(t)$ as an initial estimate of $E[S]$. If we use $\bar{W}_{L,\lambda}(t)$ as an estimate of $E[S]$ in (30), then the formula in (30) reduces to the bias approximation in (27). Thus, under these approximations, the refined estimator (28) becomes unbiased. Hence, we propose the refined estimator in (28) for light-to-moderately-loaded many-server systems with service time distributions not too far from exponential.

To better understand the consequence of non-exponential service times in the infinite-server paradigm, we now consider the $M_t/GI/\infty$ IS model with non-exponential service times. We assume that it starts empty at some time in the past (possibly in the infinite past) having bounded time-varying arrival rate $\lambda(t)$, i.i.d. service times, independent of the arrival process, with generic service-time S having cdf $G(x) \equiv P(S \leq x)$ with $E[S^2] < \infty$ and thus finite squared coefficient of variation (SCV) $c_S^2 \equiv \text{Var}(S)/E[S]^2$. Let $G^c(x) \equiv 1 - G(x)$ be the complementary cdf. Let S_e be an associated random variable with the associated stationary-excess or residual-lifetime distribution,

$$P(S_e \leq x) \equiv \frac{1}{E[S]} \int_0^x G^c(u) du \quad \text{and} \quad E[S_e^k] = \frac{E[S^{k+1}]}{(k+1)E[S]}. \quad (31)$$

For this IS model, we can characterize the conditional expected value of the remaining work $T_W^{(r)}(t)$ in (5) and (7) given $L(t)$, but it requires the full waiting-time cdf G .

THEOREM 5. (*total remaining work for the $M_t/GI/\infty$ infinite-server model*) For the $M_t/GI/\infty$ model above,

$$E[T_W^{(r)}(t)|L(t)] = \frac{L(t) \int_0^\infty \lambda(t-u)E[S-u; S>u] du}{E[L(t)]}, \quad (32)$$

for $T_W^{(r)}(t)$ in (5), where $E[S-u; S>u] = E[S-u|S>u]P(S>u)$, $E[S-u|S>u] = \int_0^\infty (G^c(x-u)/G^c(u)) dx$, and

$$E[L(t)] = \int_0^\infty \lambda(t-u)G^c(u) du = E[\lambda(t-S_e)]E[S], \quad t \geq 0. \quad (33)$$

Proof. Conditional on $L(t) = n$, the n customers remaining in service have i.i.d. service times distributed as S_t with

$$P(S_t > x) = \frac{\int_0^\infty \lambda(t-u)P(S > x+u) du}{E[L(t)]}, \quad (34)$$

for $E[L(t)]$ given in (33), by Theorem 2.1 of Goldberg and Whitt (2008), which draws on Eick et al. (1993a). There the system starts empty at time 0, but the result extends to the present setting, given that we have assumed that the arrival rate function is bounded and $E[S^2] < \infty$. The second expression in (33) is given in Theorem 1 of Eick et al. (1993a). ■

If we now invoke the BAA for the $M_t/GI/\infty$ model, then we obtain the approximation

$$E[\Delta_W(t)|\mathcal{O}_t] \approx \frac{E[T_W^{(r)}(0)|L(0)] - E[T_W^{(r)}(t)|L(t)]}{A(t)}, \quad (35)$$

where (32) can be used to compute both terms in the numerator.

In practice, we presumably would not know the full service-time cdf, so that the approximation in (35) based on Theorem 5 would not appear to be very useful, but we now show that it provides strong support for the refined estimator in (28) if the service-time is not too far from exponential. For that purpose, we observe that the complicated formula above simplifies in special cases. First, for $M_t/M/\infty$, formula (32) reduces to $E[T_W^{(r)}(t)|L(t)] = L(t)E[S]$, taking us back to (27). Second, for the stationary $M/GI/\infty$ model starting empty in the infinite past, S_t in (34) is distributed as S_e in (31), so that formula (32) reduces to $E[T_W^{(r)}(t)|L(t)] = L(t)E[S_e] = L(t)E[S](c_s^2 + 1)/2$ and (35) reduces to $E[\Delta_W(t)|\mathcal{O}_t] = (R(0) - L(t))E[S](c_s^2 + 1)/2A(t)$, depending only on the first two moments of the distribution.

This result for the stationary $M/GI/\infty$ model applies to the nonstationary $M_t/GI/\infty$ system if the arrival rate is nearly constant just prior to the two times 0 and t , where we would be applying Theorem 5. Thus, we conclude that this section provides strong support for the refined estimator $\bar{W}_{L,\lambda,r}(t)$ in (28) in the common case where (i) the arrival rate changes relatively slowly compared to the mean service time and (ii) the service-time SCV c_s^2 is not too far from 1, as is often the case in call centers, e.g., here (where $c_s^2 = 1.017$) and in Brown et al. (2005). We could obtain a further refinement if we could estimate the SCV c_s^2 .

5.4. Bias of $\bar{W}(t)$ in the Infinite-Server Paradigm

We now observe that the bias of $\bar{W}(t)$ as an estimator of W should usually not be a major factor in the infinite-server paradigm. We do so by showing that the bias is quantifiably small for an IS model. We use the $G_t/GI/\infty$ model with general, possibly nonstationary, arrival counting process A . The key assumption is that the waiting times, which coincide with the service times, are i.i.d. with mean W and independent of the arrival process. Using that independence, we can write

$$E[\bar{W}(t)|A(t) > 0] = E[E[\bar{W}(t)|A(t)]|A(t) > 0] = E[W|A(t) > 0] = W. \quad (36)$$

Given that we have defined $\bar{W}(t) \equiv 0$ when $A(t) = 0$, we have the following result.

THEOREM 6. (*conditional bias of the average waiting time in the $G_t/GI/\infty$ model*) *For the $G_t/GI/\infty$ infinite-server model, having i.i.d. service times with mean W , that are independent of a general arrival process,*

$$E[\bar{W}(t)] = WP(A(t) > 0). \quad (37)$$

For a stationary Poisson arrival process with rate λ , $W - E[\bar{W}(t)] = We^{-\lambda t}$, $t \geq 0$.

5.5. The Single-Server Paradigm

To show that the refined estimator $\bar{W}_{L,\lambda,r}(t)$ in (28) is not always good and that the bias can be analyzed exactly in some cases and can be significant, we now consider a single-server model. Let $L(t)$ be the number of customers waiting in queue in a single-server $G_t/GI/1$ queueing model with

unlimited waiting space and the first-come first-served service discipline, with a general arrival process possibly having a time-varying arrival-rate function $\lambda(t)$ and service times S_i that are independent and identically distributed (i.i.d.) and independent of the arrival process, each distributed as a random variable S having cdf $G(x)$. In addition to the model structure, we assume that we know the mean $E[S]$, which in practice may be based on a sample mean estimate.

We now assume that $T_W^{(r)}(0)$ in (5) is observable, which is reasonable because customers depart in order of arrival in the single-server model. It is also necessary for all these customers to have departed by time t , which is reasonable if t is not too small. Let $S^{(r)}(t)$ be the residual service time of the customer in service at time t , if any. In this setting, the total remaining waiting time of all customers in the system at time t is given by

$$T_W^{(r)}(t) \equiv \sum_{k=1}^{L(t)} W_k^{r,t} = L(t)S^{(r)}(t) + \sum_{k=1}^{L(t)-1} (L(t) - k)S_{k+1}, \quad (38)$$

where S_k , $k \geq 2$, are i.i.d. and independent of $L(t)$ and $S^{(r)}(t)$, but in general $L(t)$ and $S^{(r)}(t)$ are dependent. Further simplification occurs if S is exponential.

THEOREM 7. (*bias reduction for the $G_t/M/1/\infty$ model*) For the $G_t/M/1/\infty$ model,

$$E[T_W^{(r)}(t)|L(t), E[S]] = L(t)(L(t) + 1)E[S]/2. \quad (39)$$

so that, if $T_W^{(r)}(0)$ is fully observable in $[0, t]$, then

$$E[\Delta_W(t)|\mathcal{O}_t] = \frac{T_W^{(r)}(0) - L(t)(L(t) + 1)E[S]/2}{A(t)} \quad \text{for } \mathcal{O}(t) \equiv (L(t), \bar{L}(t), \bar{\lambda}(t), T_W^{(r)}(0), E[S]). \quad (40)$$

Proof. Formula (40) follows directly from (39), which in turn follows from (38) given that $S^{(r)}(t)$ has the same exponential distribution as S_1 and $1 + \dots + (n-1) = n(n-1)/2$. ■

We apply Theorem 7 to obtain the single-server refined estimator

$$\bar{W}_{L,\lambda,r,1}(t) \equiv \bar{W}_{L,\lambda}(t) - \frac{T_W^{(r)}(0) - L(t)(L(t) + 1)E[S]/2}{A(t)}. \quad (41)$$

Even if we do not know the mean $E[S]$, formulas (39)-(41) provide important insight, showing that $E[T_W^{(r)}(t)|L(t), E[S]]$ is approximately proportional to $L(t)^2$ instead of $L(t)$ as in (27) and §5.3. We next show that the bias in (40) can be significant by considering a transient $M/M/1$ example.

A Simulation Example: the $M/M/1$ Queue Starting Empty. To illustrate the bias for single-server models discussed in §5.5, we report results from a simulation experiment for the $M/M/1$ queue with mean service time $1/\mu = 1$ starting empty over the interval $[0, 10]$ for 3 values of the constant arrival rate λ : 0.7, 1.0 and 2.0. The respective 95% confidence intervals (CI's) for the exact value of $E[\bar{W}(t)]$ estimated by the sample average of 1000 replications of $\bar{W}(t)$ were 1.88 ± 0.08 , 2.70 ± 0.12 and 6.36 ± 0.19 ; the sample means are regarded as the exact values. (In an application of Little's law, these direct estimates would not be available.) To see that the refined estimator $\bar{W}_{L,\lambda,r,1}(t)$ in (41) has essentially no bias at all, without expense of wider CI's, the corresponding CI's for it based on the same 1000 replications were 1.90 ± 0.08 , 2.68 ± 0.11 and 6.38 ± 0.18 . In contrast, the unrefined $\bar{W}_{L,\lambda}(t)$ in (3) produced the corresponding tighter erroneous CI's 1.47 ± 0.06 , 1.82 ± 0.06 and 2.85 ± 0.06 . From the analysis above, we should not expect that the $G_t/M/\infty$ refined estimator (28) should perform well here. That is confirmed by the corresponding CI's 1.83 ± 0.08 , 2.46 ± 0.10 and 4.46 ± 0.11 . That is pretty good for $\lambda = 0.7$, but it misses badly for $\lambda = 2.0$.

6. Confidence Intervals for the Refined Estimator

We now see how the two statistical techniques in §§4 and 5 can be combined. We estimate confidence intervals for the refined estimators in (28) and (41) as well as the other estimators in (1) and (3).

6.1. Confidence Intervals for the Mean Wait in the Transient $M/M/1$ Queue

We now give an example in which *both* bias reduction and estimating confidence intervals contribute significantly to our understanding. To see large bias, we return to the example of the transient $M/M/1$ queue in §5.5. We now show how the sample average approach can be applied to estimate confidence intervals for the refined estimator in (41) that eliminates the bias. We now consider 10 i.i.d. samples of the same $M/M/1$ model over the interval $[0, 10]$, starting empty at time 0. We study the CI coverage by performing 1,000 replications of the entire experiment.

Table 4 shows that the unrefined estimator $\bar{W}_{L,\lambda}(t)$ in (3) does a very poor job in estimating the mean wait because of the bias, but the performance of the refined estimator $\bar{W}_{L,\lambda,r}(t)$ in (28) and the direct estimator $\bar{W}(t)$ is not too bad. It is known that residual skewness of the estimates can

degrade the performance of confidence intervals, but we find that our estimates are not extreme examples of non-normality and skewness; see §EC.4 for details. In an effort to obtain a better estimate of confidence intervals, one can consider using appropriate confidence interval inflation factor. We estimate it to be about 1.55, 1.45 and 1.05 for $\lambda = 0.7, 1.0$ and 2.0 respectively (details in §EC.4). For more discussion on skewness-adjusted CI, see Johnson (1978), Willink (2005); in context of batch means and their residual skewness and correlations, see Alexopoulos and Goldsman (2004), Tafazzoli et al. (2011), Tafazzoli and Wilson (2011) and references therein.

Table 4 Confidence intervals for the mean wait in the transient $M/M/1$ Queue for $\lambda = 0.7, 1.0$, and 2.0 . Results are based on 1000 replications of 10 i.i.d. samples of the same $M/M/1$ model over the interval $[0, 10]$, starting empty at time 0. True mean wait values are estimated using 100,000 simulation runs and assumed to be 1.8913, 2.6354 and 6.3786 for $\lambda = 0.7, 1.0$, and 2.0 , respectively.

λ	$\bar{L}(t)$	$\bar{\lambda}(t)$	$\bar{W}(t)$	cov.	$\bar{W}_{L,\lambda}(t)$	cov.	$\bar{W}_{L,\lambda,r}(t)$	cov.
0.7	1.10 ± 0.57	0.70 ± 0.17	1.89 ± 0.85	90.3%	1.46 ± 0.57	58.3%	1.88 ± 0.82	89.9%
1.0	1.91 ± 0.88	1.00 ± 0.21	2.63 ± 1.16	90.5%	1.80 ± 0.63	31.5%	2.62 ± 1.11	92.2%
2.0	5.82 ± 1.74	2.00 ± 0.29	6.36 ± 2.07	91.3%	2.83 ± 0.63	0.0%	6.38 ± 1.95	93.1%

6.2. Evaluating the Refined Estimator with the Call Center Data

Given that the call center should approximately fit the infinite-server paradigm and that the waiting-time distribution is approximately exponential, we can apply equation (29) to see that the bias should be relatively small in the call center example. We now use data from the 18 weekdays in May 2001 for the call center example in §3 to confirm that observation and show that the refined estimator in (28) is effective in reducing the bias.

Since we observe strong day-to-day variation in the average waiting times, we do not try to estimate the overall mean over all days, but aim to estimate the mean of specified intervals on each day (for sample averages over all days and their associated confidence interval, see Section 6.3). In particular, we compute the average over the 18 days of the absolute errors $|\bar{W}_{L,\lambda}(t) - \bar{W}(t)|$ (AAE) and associated average squared errors (ASE) for each of the 17 hours and 34 half hours of the day. We choose hours and half hours, because they represent typical staffing intervals in call centers; see

Green et al. (2007). Table 5 highlights the results; AAE and ASE of each subinterval (hours and half hours) are again averaged over the intervals [6, 10], [10, 16], [16, 23] and all day. More details appear in Tables 15 and 16 in Kim and Whitt (2012)

Table 5 Comparison of the refined estimator $\bar{W}_{L,\lambda,r}(t)$ in (28) to the unrefined estimator $\bar{W}_{L,\lambda}(t)$ in (3): Average over the day of the average absolute errors (AAE) and average squared errors (ASE) for each time interval over 18 weekdays in the call center example.

Subinterval Length	Intervals averaged over	unrefined in (3)			refined in (28)		
		$\bar{W}_{L,\lambda}(t)$	AAE	ASE	$\bar{W}_{L,\lambda,r}(t)$	AAE	ASE
hours	[6, 10]	3.32	0.241	0.117	3.54	0.082	0.018
	[10, 16]	3.61	0.076	0.010	3.60	0.058	0.006
	[16, 23]	4.46	0.271	0.160	4.28	0.153	0.057
	all	3.89	0.195	0.097	3.86	0.103	0.030
half hours	[6, 10]	3.27	0.303	0.198	3.49	0.169	0.068
	[10, 16]	3.62	0.161	0.052	3.60	0.110	0.020
	[16, 23]	4.55	0.533	0.673	4.25	0.340	0.322
	all	3.92	0.347	0.342	3.84	0.219	0.156

Table 5 shows that the refined estimator reduces the AAE from 0.195 (about 5.0% of the overall average wait, 3.89) to 0.103 (2.6%) for hours over all hours, while the refined estimator reduces the AAE from 0.347 (8.9%) to 0.219 (5.6%) for half hours over all half hours. In both cases, there is more bias and more bias reduction at the ends of the day when the system is nonstationary. In addition, we note that the unrefined estimator underestimates $\bar{W}(t)$ during [6, 10] when the arrival rate is increasing, and that it overestimates $\bar{W}(t)$ during [16, 23] when the arrival rate is decreasing, as expected.

6.3. Sample Averages Over Separate Days

For many service systems, whether stationary or not, we may be able to estimate CI's for $E[\bar{W}(t)]$ in (1) without observing the waiting times via $E[\bar{W}_{L,\lambda}(t)]$ in (3) using sample averages over multiple days, regarding those days as approximately i.i.d. We assume that the time average operation makes the vector $(\bar{L}(t), \bar{\lambda}(t))$ approximately Gaussian for each day. Thus, by Theorem 4, the associated random variable $\bar{W}_{L,\lambda}(t)$ should be approximately Gaussian as well with (unknown) variance given in (15). We also assume that any refinement $\bar{W}_{L,\lambda,r}(t)$ is approximately Gaussian as well.

Based on n days regarded as i.i.d., we can construct CI in the usual way. Let X_i denote the time average $\bar{W}_{L,\lambda}(t)$ or (preferably) its refinement $\bar{W}_{L,\lambda,r}(t)$ based on the bias analysis described in §5 for day i . Let the sample mean and variance be

$$\bar{X}_n \equiv \frac{1}{n} \sum_{i=1}^n X_i \quad \text{and} \quad S_n^2 \equiv \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X}_n)^2. \quad (42)$$

Then $(\bar{X}_n - E[\bar{W}(t)])/\sqrt{S_n^2/n}$ should be approximately distributed as t_{n-1} , Student t with $n-1$ degrees of freedom. Then $\bar{X}_n \pm t_{\alpha/2, n-1} S_n/\sqrt{n}$ is a $1-\alpha$ CI for $E[\bar{W}(t)]$.

To assess how well indirect estimators perform in estimating $E[\bar{W}(t)]$ over separate days and in different settings, we again consider our call center data and divide each day into 3 intervals, $[6, 10]$, $[10, 16]$ and $[16, 23]$ so that the arrival rate is increasing in $[6, 10]$, approximately stationary in $[10, 16]$ and decreasing in $[16, 23]$. The performance of two indirect estimators, the refined estimator $\bar{W}_{L,\lambda,r}(t)$ in (28) and the unrefined estimator $\bar{W}_{L,\lambda}(t)$ in (3), as well as that of direct estimator is illustrated in Table 6. (Additional estimation results appear in Tables 17-19 of Kim and Whitt (2012).) We see that the refined estimator $\bar{W}_{L,\lambda,r}(t)$ behaves very similar to the direct estimator in all cases. The unrefined estimator performs well in the stationary region $[10, 16]$, but shows the impact of bias in nonstationary regions, $[6, 10]$ and $[16, 23]$, as expected.

Table 6 Estimating $E[\bar{W}(t)]$ and its associated 95% confidence interval over 18 weekdays in the call center example: comparison of the refined estimator $\bar{W}_{L,\lambda,r}(t)$ in (28) to the unrefined estimator $\bar{W}_{L,\lambda}(t)$ in (3)

Intervals	direct estimator $\bar{W}(t)$	unrefined in (3) $\bar{W}_{L,\lambda}(t)$	refined in (28) $\bar{W}_{L,\lambda,r}(t)$
$[6, 10]$	3.47 ± 0.22	3.35 ± 0.23	3.47 ± 0.23
$[10, 16]$	3.60 ± 0.11	3.61 ± 0.11	3.60 ± 0.11
$[16, 23]$	4.24 ± 0.26	4.35 ± 0.26	4.22 ± 0.25

7. Conclusions

Little's law is an important theoretical cornerstone of operations research, but it does not apply directly to applications involving measurements over finite time intervals. As reviewed in §2.3, it is possible to modify the definitions so that the relation $\bar{L}(t) = \bar{\lambda}(t)\bar{W}(t)$ always holds for finite averages, but we advocate not doing so. Instead, we advocate taking a statistical approach, estimating

confidence intervals (§4) and considering modified estimators that reduce bias (§5), which exploit the extended finite-time Little's law in Theorem 2. We have illustrated the statistical approach by applying it to the call center example in §3. We have focused on the problem of estimating the unknown mean values W and $E[\bar{W}(t)]$ using $\bar{W}_{L,\lambda}(t) \equiv \bar{L}(t)/\bar{\lambda}(t)$ when the waiting times cannot be directly observed.

Acknowledgements. We thank Avishai Mandelbaum, Galit Yom-Tov, Ella Nadjharov and the Center for Service Enterprise Engineering (SEE) at the Technion for access to the SEE call center data and advice about its use. We thank the Samsung Foundation and NSF for support (NSF grant CMMI 1066372).

References

- Alexopoulos, C., D. Goldsman. 2004. To batch or not to batch. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* **14** 76–114.
- Alexopoulos, C., N. T. Argon, D. Goldsman, N. M. Steiger, A. Tafazzoli, J. R. Wilson. 2007. Efficient computation of overlapping variance estimators for simulation. *INFORMS J. on Computing* **19** 314–327.
- Asmussen, S. 2003. *Applied Probability and Queues*, second ed., Springer, New York.
- Asmussen, S., P. W. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*, Springer, New York.
- Baccelli, F., P. Bremaud. 2003. *Elements of Queueing Theory: Palm-Martingale Calculus and Stochastic Recurrences*, second ed., Springer, New York.
- Brown, L., N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn, L. Zhao. 2005. Statistical analysis of a telephone call center: a queueing-science perspective. *J. Amer. Statist. Assoc.* **100** 36–50.
- P. J. Brockwell, R. A. Davis. 1991. *Time Series Theory and Methods*, second ed., Springer, New York.
- Buzen, J. P. 1976. Fundamental operational laws of computer system performance. *Acta Informatica* **7** 167–182.
- Conway, R. W. 1963. Some tactical problems in digital simulation. *Management Sci.* **10** 47–61.
- Denning, P. J., Buzen, J. P. 1978. The operational analysis of queueing network models. *Computing Surveys* **10** 225–261.

- Eick, S. G., W. A. Massey, W. Whitt. 1993a. The physics of The $M_t/G/\infty$ queue. *Operations Research* **41** 731–742.
- Eick, S. G., W. A. Massey, W. Whitt. 1993b. $M_t/G/\infty$ queues with sinusoidal arrival rates. *Management Sci.* **39**, 241–252.
- El-Taha, M., S. Stidham, Jr. 1999. *Sample-Path Analysis of Queueing Systems*, Kluwer, Boston.
- Glynn, P. W., W. Whitt. 1986. A central-limit-theorem version of $L = \lambda W$. *Queueing Systems* **1** 191–215.
- Glynn, P. W., W. Whitt. 1987. Sufficient conditions for functional-limit-theorem versions of $L = \lambda W$. *Queueing Systems* **1** 279–287.
- Glynn, P. W., W. Whitt. 1989. Indirect estimation via $L = \lambda W$. *Oper. Res.* **37** 82–103.
- Goldberg, D. A., W. Whitt. 2008. The last departure time from an $M_t/G/\infty$ queue with a terminating arrival process. *Queueing Systems* **58** 77–104.
- Green, L. V., P. J. Kolesar, W. Whitt. 2007. Coping with time-varying demand when setting staffing requirements for a service system. *Production and Operations Management* **16** 13–39.
- Jennings, O. B., A. Madelbaum, W. A. Massey, W. Whitt. 1996. Server staffing to meet time-varying demand. *Management Sci.* **42** 1383–1394.
- Jewell, W. S. 1967. A simple proof of $L = \lambda W$. *Oper Res.* **15** 1109–1116.
- Johnson, N. J. 1978. Modified t tests and confidence intervals for asymmetrical populations. *Journal of the American Statistical Association* **73** 536–544.
- Kim, S.-H., W. Whitt. 2012. Statistical analysis with Little's law, supplementary material: Technical report. Available from: <http://www.columbia.edu/~ww2040/allpapers.htm>
- Little, J. D. C. 1961. A proof of the queueing formula: $L = \lambda W$. *Oper. Res.* **9** 383–387.
- Little, J. D. C. 2011. Little's law as viewed on its 50th anniversary. *Oper. Res.* **59** 536–539.
- Little, J. D. C., S. C. Graves. 2008. Little's law. Ch. 5 in *Building Intuition: Insights from Basic Operations Management Models and Principles*, D. Chhajed, T. J. Lowe (eds.), Springer, New York, 81–100.
- Lovejoy, W. S., J. S. Desmond. 2011. Little's law flow analysis of observation unit impact and sizing. *Acad. Emergency Medicine* **18** 183–189.
- Mandelbaum, A. 2011. Little's law over a finite horizon. Pages 17.1–17.6 in Teaching

- notes on Little's law in a course on Service Engineering, October 2011. Available at: <http://iew3.technion.ac.il/serveng/Lectures/lectures.html> (Accessed August 3, 2012)
- Mandelbaum, A. 2012. Service Engineering of Stochastic Networks web page: <http://iew3.technion.ac.il/serveng/>
- Serfozo, R. 1999. *Introduction to Stochastic Networks*, Springer, New York.
- Sigman, K. 1995. *Stationary Marked Point Processes, An Intuitive Approach*, Chapman and Hall, New York.
- Stidham, S., Jr. 1974. A last word on $L = \lambda W$. *Oper. Res.* **22** 417–421.
- Tafazzoli, A., N. M. Steiger, J. R. Wilson. 2011. N-Skart: A nonsequential skewness- and autoregression-adjusted batch-means procedure for simulation analysis *IEEE Trans. Aut. Control* **56** 254–264.
- Tafazzoli, A., J. R. Wilson. 2011. Skart: A skewness- and autoregression-adjusted batch-means procedure for simulation analysis *IIE Trans.* **43** 110–128.
- Whitt, W. 1991. A review of $L = \lambda W$ and extensions. *Queueing Systems* **9** 235–268.
- Whitt, W. 2012. Extending the FCLT version of $L = \lambda W$. *Oper. Res. Letters* **40** 230–234.
- Willink, R. 2005. A confidence interval and test for the mean of an asymmetric distribution. *Communications in Statistics Theory and Methods* **34** 753–766.

This page is intentionally blank. Proper e-companion title page, with INFORMS branding and exact metadata of the main paper, will be produced by the INFORMS office when the issue is being assembled.

E-Companion

EC.1. Overview

In this e-companion we supplement the main paper. In §EC.2 we provide the stationary test results for our call center data. In §EC.3 we supplement §3 and §5.3 by showing how idealized simulation models for the call center example were constructed. In §EC.4 we elaborate on the confidence intervals for the mean wait in the transient $M/M/1$ queue discussed in §6.1. We show that the distribution of the estimates are not too far from being normal, and show ways to adjust the confidence interval halfwidths to get targeted coverage.

Additional material is in a report available online from the authors' web pages (Kim and Whitt 2012). Following an introduction in §1, in §2 we present additional information about the call center data used. In §3 we discuss an alternative way to construct the confidence intervals using the method of batch means and in §4 we elaborate on the bias discussed in §5, discussing the bias in the estimator $\bar{W}(t)$ in (1) for W in a stationary setting in §4.1. In §4.2 we discuss the bias in the alternative estimator $\bar{W}_{L,\lambda}(t)$ in (3) for $E[W(t)]$ in a nonstationary setting and W in a stationary setting. In §5 we introduce an alternative algorithm to estimate confidence intervals in approximately stationary intervals by batch means, exploiting Theorem 2. However, this algorithm did not improve the estimation for the call center example, so we do not emphasize it. The negative result itself is interesting, because it is natural to consider such alternatives. It remains to be seen if the new algorithm can be useful in other contexts. Finally, in §6 we present additional figures and tables.

EC.2. Results for the Stationarity Tests

In this section, we provide detailed results for the three tests for stationarity applied to the call center data in §3.2. The three tests are the turning points test, the difference-sign test and the rank test for randomness, all discussed on p. 312 of Brockwell and Davis (1991). Given a sequence of observations $\{y_1, \dots, y_n\}$, these tests basically check the hypothesis that $\{y_i\}$ is an i.i.d. sequence.

For *the turning points test*, we say that the data has a turning point at time i if $y_{i-1} < y_i$ and $y_i > y_{i+1}$ or $y_{i-1} > y_i$ and $y_i < y_{i+1}$. Let T be the number of turning points of the sequence $\{y_i\}$. If $\{y_i\}$ are observations of an i.i.d. sequence, the expected number of turning points is $\mu_T = 2(n-2)/3$ with variance $\sigma_T^2 = (16n-29)/90$, and T is approximately $N(\mu_T, \sigma_T^2)$. A large value of $T - \mu_T$ means the sequence is fluctuating more rapidly than expected for a random sequence. A large (in magnitude) negative value of $T - \mu_T$ suggests a positive correlation between neighboring observations. *The difference-sign test* counts the number (S) of values of i such that $y_i > y_{i-1}$, $i = 2, \dots, n$. Under the i.i.d. sequence assumption, $\mu_S = \frac{1}{2}(n-1)$ with variance $\sigma_S^2 = (n+1)/12$, and S is approximately $N(\mu_S, \sigma_S^2)$. *The rank test* counts the number (P) of pairs (i, j) such that $y_j > y_i$, $j > i$, $i = 1, \dots, n-1$. If $\{y_i\}$ are observations of an i.i.d. sequence, then $\mu_P = \frac{1}{4}n(n-1)$ with variance $\sigma_P^2 = n(n-1)(2n+5)/8$, and P is approximately $N(\mu_P, \sigma_P^2)$. A large positive (negative) value of $S - \mu_S$ or $P - \mu_P$ indicates an increasing (decreasing) trend in $\{y_i\}$.

We apply the aforementioned three tests on arrival counts over successive subintervals of varying lengths, 1, 5 and 10 minutes. Table EC.1 provides the result. If we use significance level $\alpha = 0.95$, we have $\Phi_{1-\alpha/2} = 1.96$. The test results depend on the subinterval length, but they mainly accept the hypothesis that the arrival counts are i.i.d sequence in $[10, 16]$, but reject it outside of the interval $[10, 16]$. The results of the rank test, which is particularly useful for detecting a linear trend, is especially strong.

EC.3. Details for the Simulation Experiments in Section 4

In §1.2.2, we suggest applying simulation to study how the estimation procedures proposed in this paper work for an idealized queueing model of the system. In this section, we describe in detail how we construct the idealized simulation models in §3.2 for the call center example.

EC.3.1. Fitted Arrival Rate Function

In order to construct an idealized model to evaluate how the estimation procedure perform for the actual data, we need to mimic the behavior of the actual arrival process of the call center as much as possible. Given arrival rates measured in units of 10 minutes over the working day $[6, 23]$, we fit

Table EC.1 Testing for stationarity: the assumption that $\{y_i\}$ are observations from an i.i.d sequence is rejected if $|\cdot - \mu_\cdot|/\sigma_\cdot > \Phi_{1-\alpha/2}$ where $\Phi_{1-\alpha/2}$ is the $1 - \alpha/2$ percentage point of a standard normal distribution.

Subinterval	Interval	T	μ_t	$ T - \mu_T /\sigma_T$	S	μ_S	$ S - \mu_S /\sigma_S$	P	μ_P	$ P - \mu_P /\sigma_P$
1	[6, 10]	112	158.7	7.17	96	119.5	5.24	21517	14340.0	3.85
	[10, 16]	207	238.7	3.97	159	179.5	3.74	29207	32310.0	0.91
	[16, 23]	179	278.7	11.56	144	209.5	11.06	15294	43995.0	6.66
5	[6, 10]	21	30.7	3.37	25	23.5	0.74	968	564.0	2.39
	[10, 16]	49	46.7	0.66	31	35.5	1.82	1208	1278.0	0.23
	[16, 23]	44	54.7	2.79	36	41.5	2.07	567	1743.0	3.03
10	[6, 10]	12	14.7	1.34	14	11.5	1.73	244	138.0	1.75
	[10, 16]	19	22.7	1.49	18	17.5	0.28	299	315.0	0.15
	[16, 23]	24	26.7	1.00	16	20.5	2.38	109	430.5	2.32

the arrival rate function to a continuous piecewise-linear function, with one increasing piece over $[6, 10]$, a constant piece over $[10, 16]$ and two decreasing linear pieces over $[16, 18]$ and $[18, 23]$. We force two extra constraints: (i) the arrival rate starts and ends at 0, and hence the arrival rate is 0 at $t = 6$ and $t = 23$ and (ii) the arrival rate of successive pieces agree at endpoints. Figure EC.1 illustrates the result. The exact arrival rate function is given by

$$\lambda(t) = \begin{cases} 140(t - 6) & \text{on } [6, 10], \\ 560 & \text{on } [10, 16], \\ 560 - 230(t - 16) & \text{on } [16, 18], \\ 100 - 20(t - 18) & \text{on } [18, 23]. \end{cases} \quad (\text{EC.1})$$

We used the same 1000 arrival sample paths for all the infinite-server and finite-server models used in this paper. We generated these arrival processes by thinning a homogeneous arrival process with rate $\lambda^* = 560$. The homogeneous Poisson process generates potential arrivals. We then let a potential arrival at time t be an actual arrival in the nonhomogeneous arrival process with probability $\lambda(t)/\lambda^*$.

EC.3.2. Histogram of the Waiting Time Distribution in the Call Center

We assumed that all the service times were i.i.d. with a distribution obtained to match the observed waiting time distribution. In particular, we consider the waiting times (time spent in the system) over the interval $[10, 16]$ on May 25, 2001. Figure EC.2 shows a histogram of the waiting times.

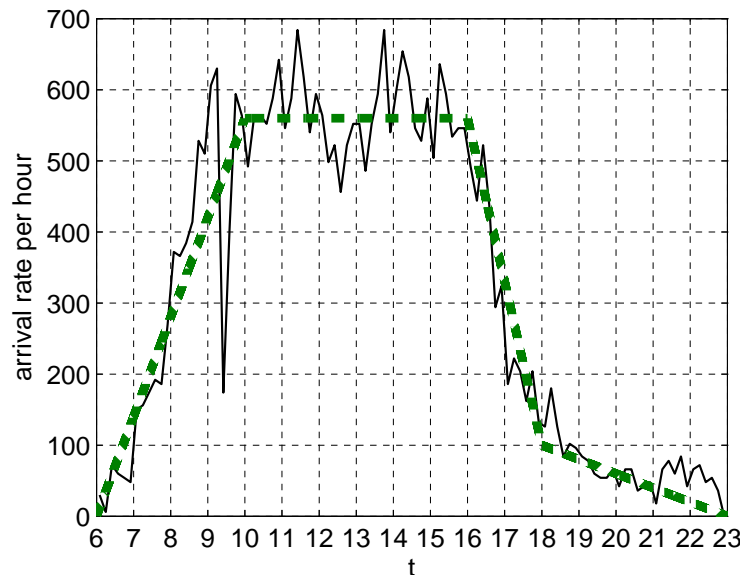


Figure EC.1 Fitted arrival rate function for the arrivals at the call center on May 25, 2001.

An exponential approximation with mean 3.38 was found to be a good fit and was used. Thus, service times were generated according to an exponential distribution with mean 3.38 for 1,000 replications. The same set of generated service times were used for all the infinite-server and finite-server models used in this paper. (Consistent with many call center empirical studies Brown et al. (2005), a lognormal distribution with mean 3.38 and squared coefficient of variation $c^2 = 1.017$ was also found to be a good fit, even better in the neighborhood of the origin, but we were not concerned that the mode falls to the right of 0. Simulation shows that the results are not significantly altered by using the fitted lognormal distribution. The exponential distribution makes the staffing easier for the $M_t/GI/s_t$ model; e.g., by applying Eick et al. (1993a).)

EC.3.3. Staffing for the $M_t/GI/s_t$ Models

For the call center example, we have data on the arrival times and waiting times as well as the number in system $L(s)$, $0 \leq s \leq t$, but we do not have data on the staffing and the complex call routing. Thus, in order to evaluate the estimation procedures, we simulate the single-class single-service-pool $M_t/GI/\infty$ IS model and associated $M_t/GI/s_t$ models with time-varying staffing levels chosen to yield good performance. Specifically, we use the fitted arrival rate function (from

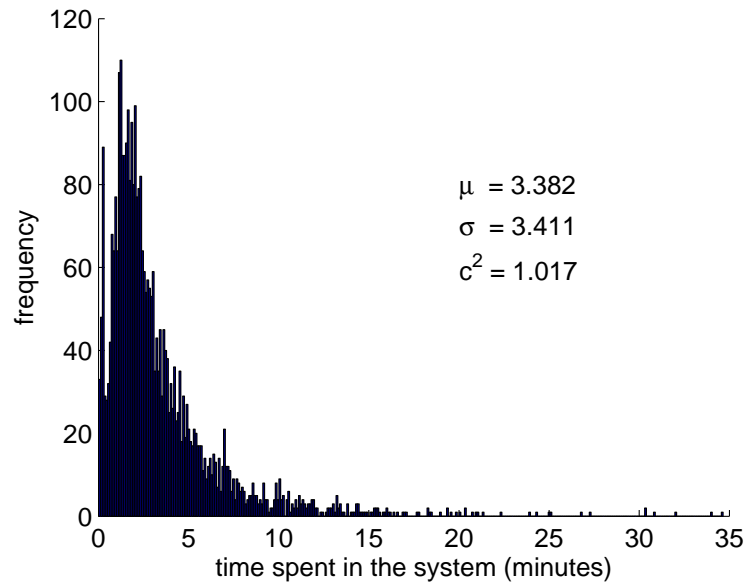


Figure EC.2 The histogram (empirical distribution) of the times spent in the system of all arrivals during the interval $[10, 16]$ on May 25, 2001.

§EC.3.1) and assume service times are exponentially distributed with mean 3.38 (from §EC.3.2). We then compute $m(t)$, the *offered load* which is the time-varying mean number of busy servers in the IS model, using formulas (6) and (7) of Jennings et al. (1996). Finally, the staffing function $s(t)$ is determined by the SRS formula using a range of Quality-of-Service (QoS) parameters, $\beta = 2.5, 2.0, 1.5, 1.0$. Figure EC.3 illustrates the offered load as well as staffing levels for different values of β .

EC.3.4. Sample Paths for Different Values of β

One way to diagnose whether simulated $M_t/GI/s_t$ models with different values of β are working as expected is to plot some sample paths of the number in the system, $L(t)$, along with the staffing levels, $s(t)$. In Figures EC.4 - EC.7, we provide one single sample path of each $M_t/GI/s_t$ model with different values of $\beta = 2.5, 2.0, 1.5, 1.0$ to illustrate that our simulation models work as expected.

EC.4. More on Confidence Intervals for the $M/M/1$ Example

In this section, we provide more details on confidence intervals for the mean wait in the transient $M/M/1$ queue, discussed in §6.1. Recall that we have 10 i.i.d. samples of the same $M/M/1$ model

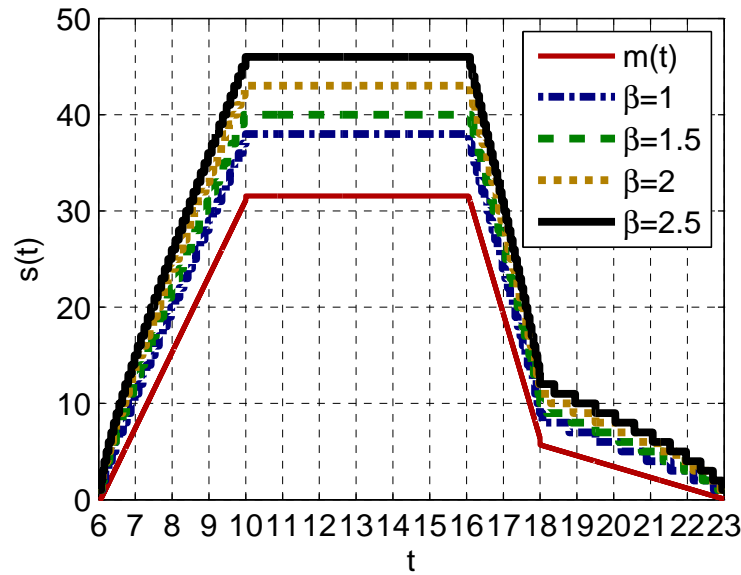


Figure EC.3 The offered load, $m(t)$, and staffing levels, $s(t)$, for different values of β .

over the interval $[0, 10]$, starting empty at time 0. To see how the sample average approach in §6.3 can be applied to estimate CIs for the refined estimator in (41), CI coverage was studied by performing 1,000 replications of the entire experiment.

In Table 4, we observed that the unrefined estimator $\bar{W}_{L,\lambda}(t)$ in (3) does a very poor job in estimating the mean wait because of the bias. The performance of the refined estimator $\bar{W}_{L,\lambda,r}(t)$ in (28) and the direct estimator $\bar{W}(t)$ is not perfect, either, yielding coverage of about 90% instead of the targeted 95%. To see whether this coverage issue is due to the residual skewness of the estimates, we perform the Shapiro-Francia normality test, discussed on p. 314 of Brockwell and Davis (1991). The Shapiro-Francia normality test is specifically designed to address issues related with large sample sizes. Table EC.2 gives the test results; normality is rejected at 5% significance level in most cases. We also plot the histogram of the estimators as well as their normality plots, from which we can see evidence of the heavy tails (skewness). However, we see that these are not extreme examples of non-normality and skewness.

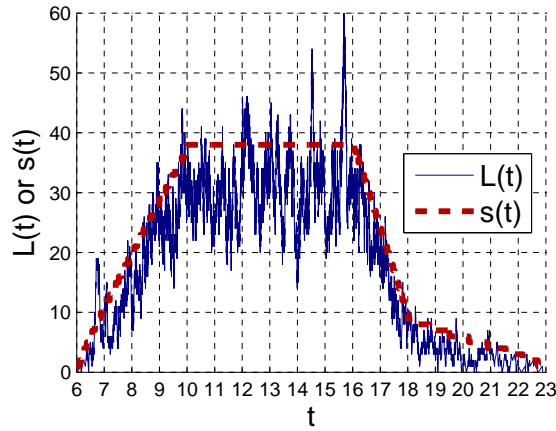


Figure EC.4 One sample path of $M_t/GI/s_t$ with $\beta = 1.0$.

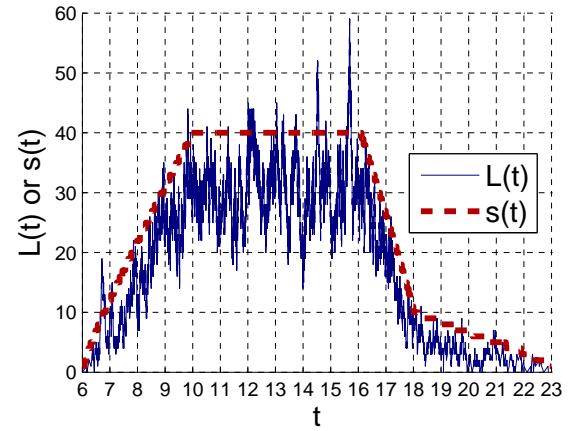


Figure EC.5 One sample path of $M_t/GI/s_t$ with $\beta = 1.5$.

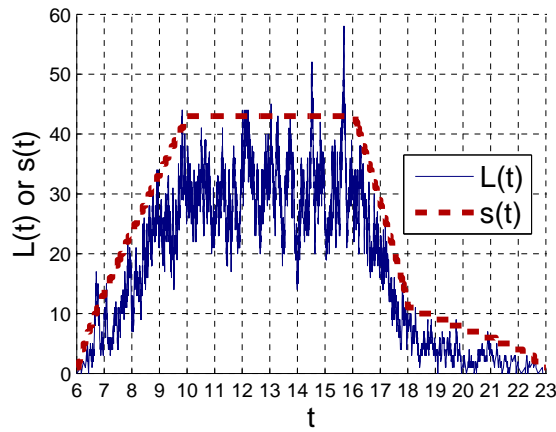


Figure EC.6 One sample path of $M_t/GI/s_t$ with $\beta = 2.0$.

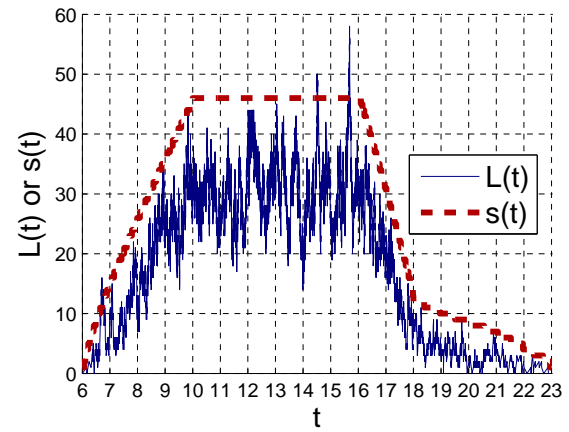


Figure EC.7 One sample path of $M_t/GI/s_t$ with $\beta = 2.5$.

Table EC.2 Shapiro-Francia normality test (discussed on p. 314 of Brockwell and Davis (1991)) of the direct and indirect mean waiting time estimator values over 1000 replications.

	$\bar{W}(t)$		$\bar{W}_{L,\lambda}(t)$		$\bar{W}_{L,\lambda,r}(t)$	
λ	R^2	p-value	R^2	p-value	R^2	p-value
0.7	0.9794	0.0000	0.9941	0.0011	0.9813	0.0000
1.0	0.9892	0.0000	0.9936	0.0006	0.9905	0.0000
2.0	0.9945	0.0019	0.9990	0.8284	0.9977	0.1562

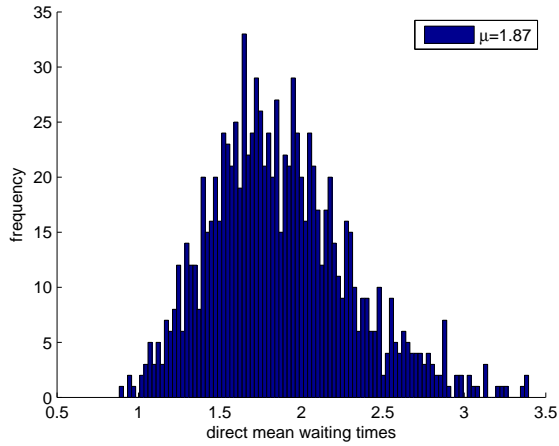


Figure EC.8 Histogram of $\bar{W}(t)$ for the $M/M/1$ model with $\lambda = 0.7$.

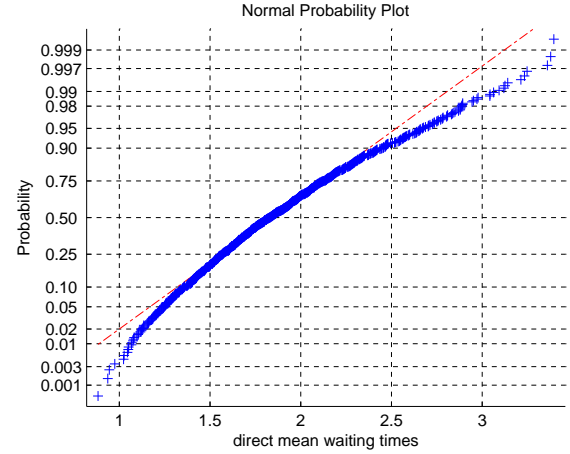


Figure EC.9 Normality plot of $\bar{W}(t)$ for the $M/M/1$ model with $\lambda = 0.7$.

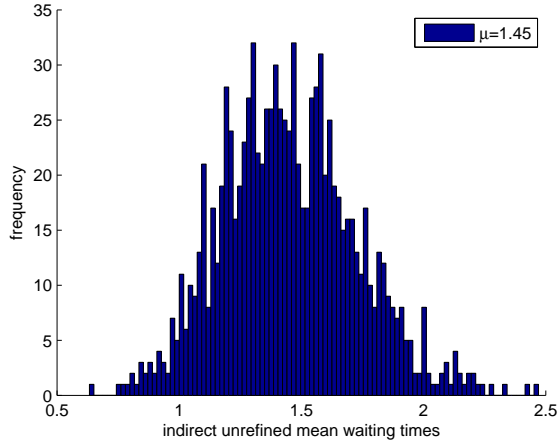


Figure EC.10 Histogram of $\bar{W}_{L,\lambda}(t)$ for the $M/M/1$ model with $\lambda = 0.7$.

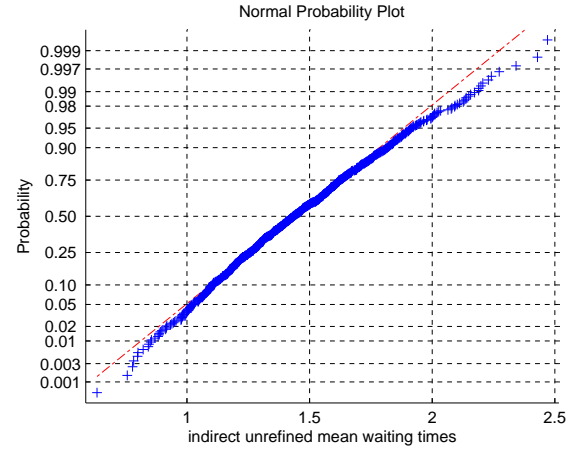


Figure EC.11 Normality plot of $\bar{W}_{L,\lambda}(t)$ for the $M/M/1$ model with $\lambda = 0.7$.

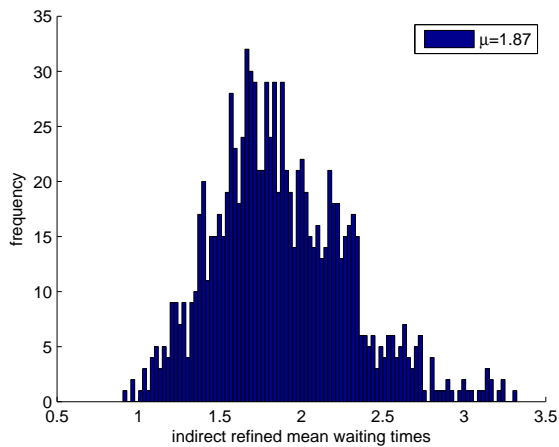


Figure EC.12 Histogram of $\bar{W}_{L,\lambda,r}(t)$ for the $M/M/1$ model with $\lambda = 0.7$.

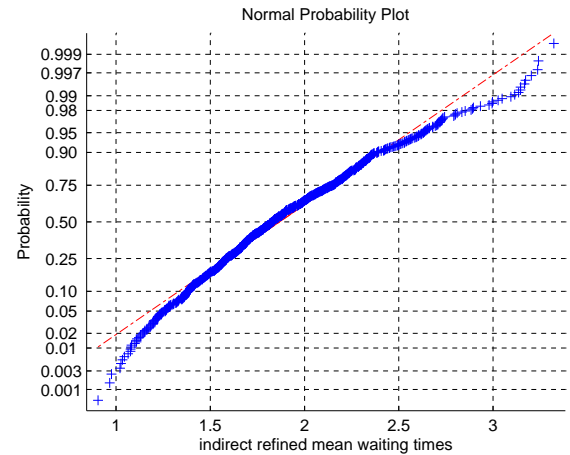


Figure EC.13 Normality plot of $\bar{W}_{L,\lambda,r}(t)$ for the $M/M/1$ model with $\lambda = 0.7$.

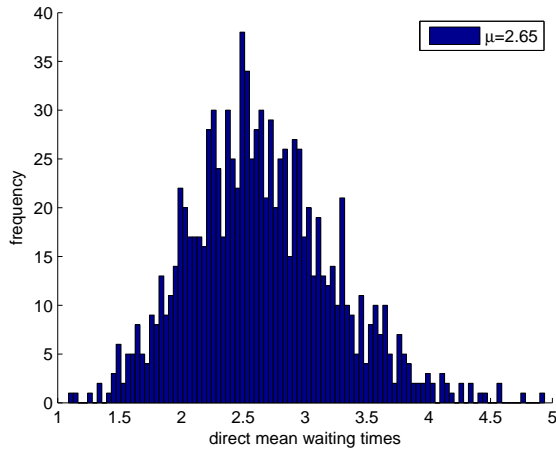


Figure EC.14 Histogram of $\bar{W}(t)$ for the $M/M/1$ model with $\lambda = 1.0$.

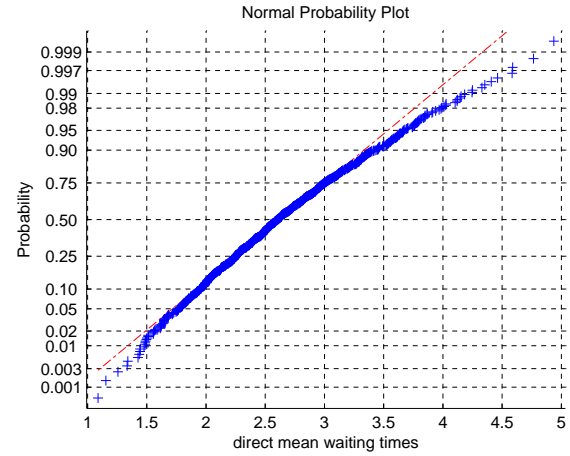


Figure EC.15 Normality plot of $\bar{W}(t)$ for the $M/M/1$ model with $\lambda = 1.0$.

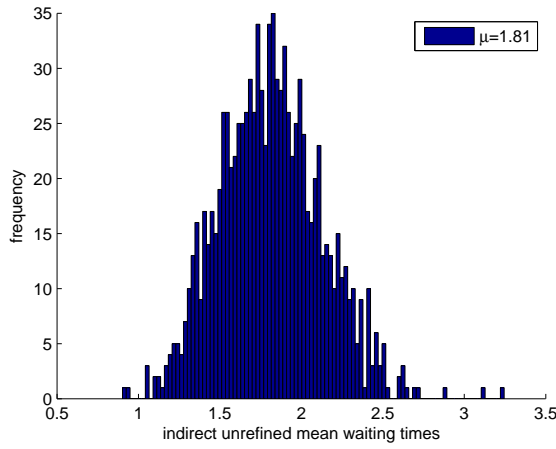


Figure EC.16 Histogram of $\bar{W}_{L,\lambda}(t)$ for the $M/M/1$ model with $\lambda = 1.0$.

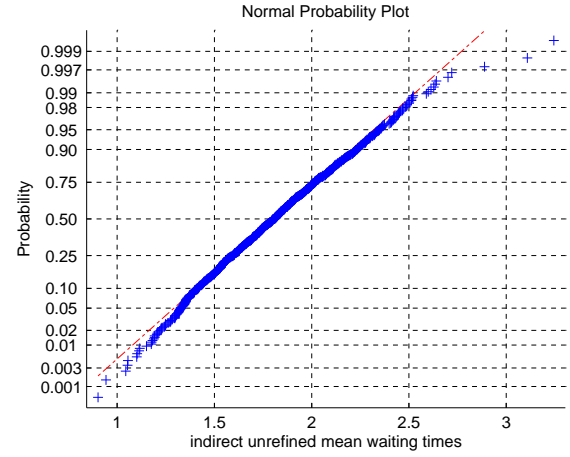


Figure EC.17 Normality plot of $\bar{W}_{L,\lambda}(t)$ for the $M/M/1$ model with $\lambda = 1.0$.

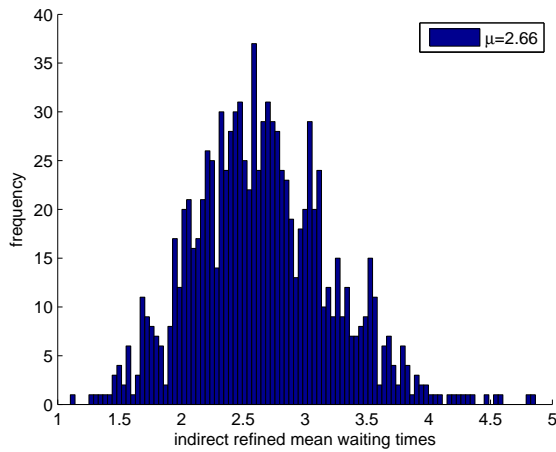


Figure EC.18 Histogram of $\bar{W}_{L,\lambda,r}(t)$ for the $M/M/1$ model with $\lambda = 1.0$.

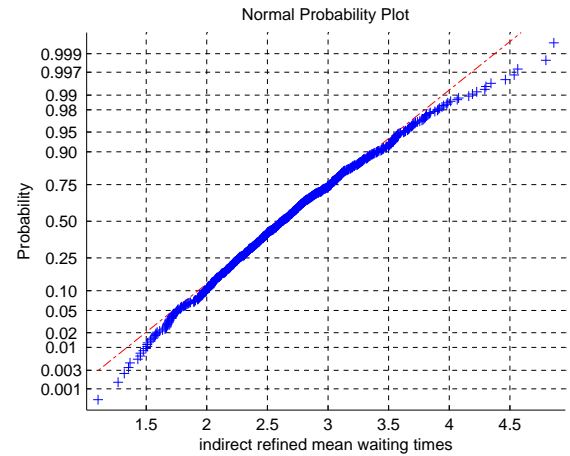


Figure EC.19 Normality plot of $\bar{W}_{L,\lambda,r}(t)$ for the $M/M/1$ model with $\lambda = 1.0$.

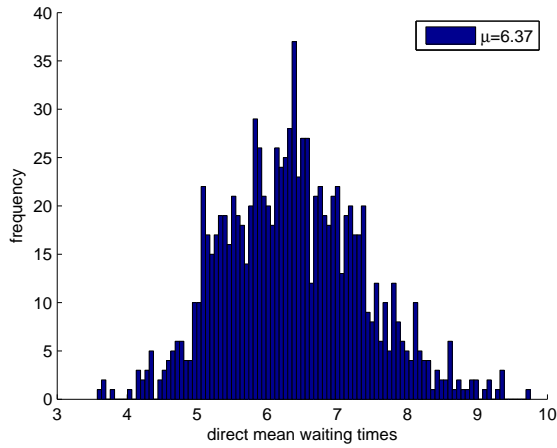


Figure EC.20 Histogram of $\bar{W}(t)$ for the $M/M/1$ model with $\lambda = 2.0$.

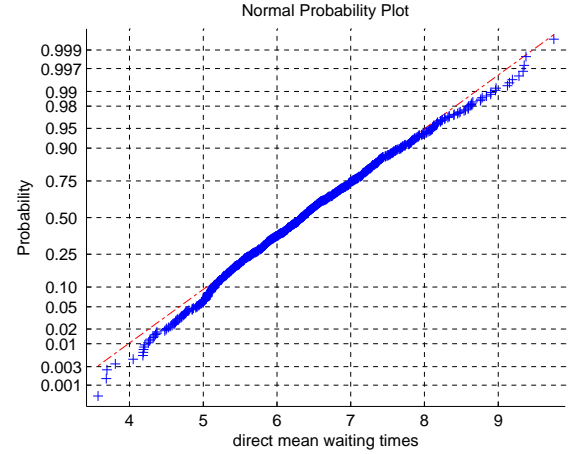


Figure EC.21 Normality plot of $\bar{W}(t)$ for the $M/M/1$ model with $\lambda = 2.0$.

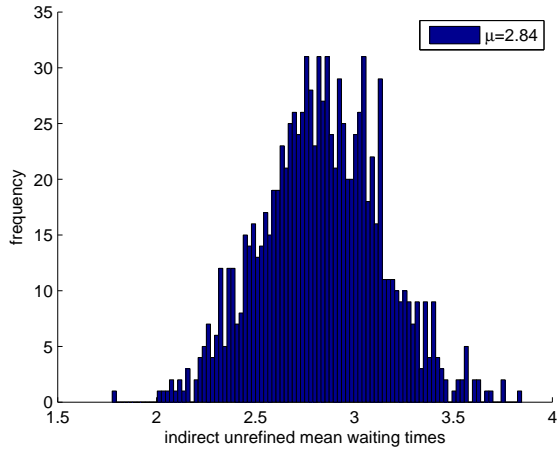


Figure EC.22 Histogram of $\bar{W}_{L,\lambda}(t)$ for the $M/M/1$ model with $\lambda = 2.0$.

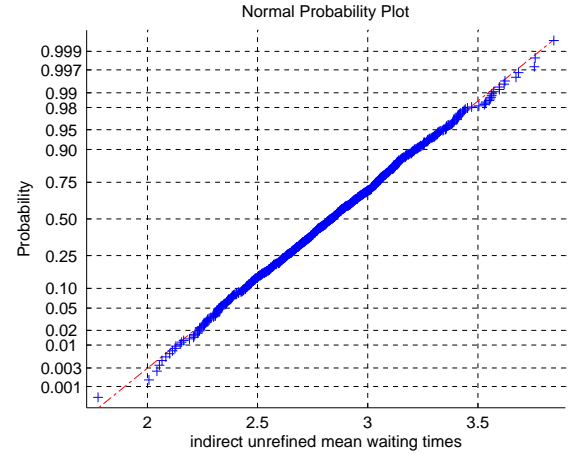


Figure EC.23 Normality plot of $\bar{W}_{L,\lambda}(t)$ for the $M/M/1$ model with $\lambda = 2.0$.

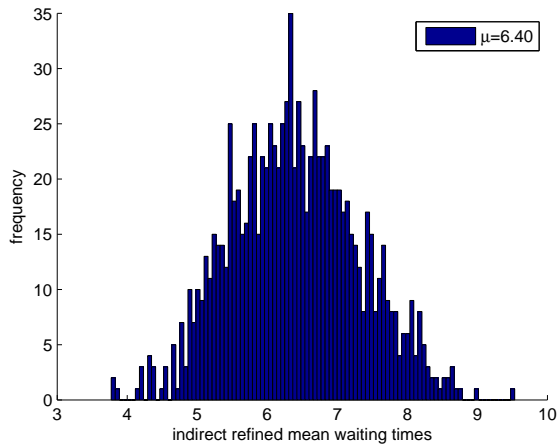


Figure EC.24 Histogram of $\bar{W}_{L,\lambda,r}(t)$ for the $M/M/1$ model with $\lambda = 2.0$.

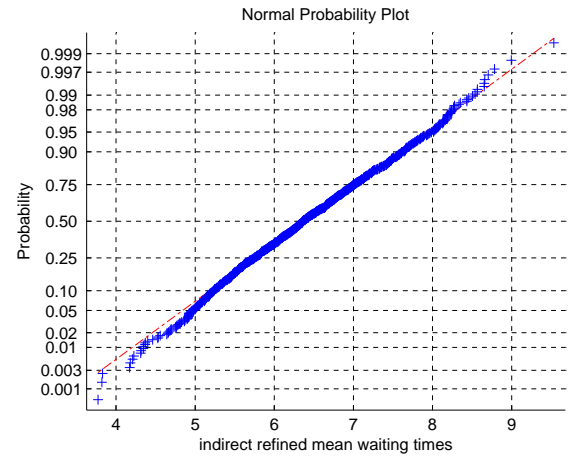


Figure EC.25 Normality plot of $\bar{W}_{L,\lambda,r}(t)$ for the $M/M/1$ model with $\lambda = 2.0$.

One way to try to obtain a better estimate of confidence interval is to use appropriate confidence interval inflation factor. The idea is to estimate the inflation factor x such that mean $\pm x$ confidence interval halfwidth has targeted (e.g., 95%) coverage. For our transient $M/M/1$ model, we test inflation factors ranging from 1.00 to 2.00, with increments of size 0.05. Table EC.3 provides the results for different values of the inflation factor. Given this result, we would estimate it to be about 1.55, 1.45 and 1.05 for $\lambda = 0.7$, 1.0 and 2.0, respectively.

Table EC.3 Estimating confidence interval inflation factor.

	$\lambda = 0.7$		$\lambda = 1.0$		$\lambda = 2.0$	
x	Cov. of $\bar{W}(t)$	Cov. of $\bar{W}_{L,\lambda,r}(t)$	Cov. of $\bar{W}(t)$	Cov. of $\bar{W}_{L,\lambda,r}(t)$	Cov. of $\bar{W}(t)$	Cov. of $\bar{W}_{L,\lambda,r}(t)$
1.00	87.3%	87.8%	90.2%	90.7%	94.0%	95.1%
1.05	88.5%	89.3%	91.4%	91.9%	95.4%	95.8%
1.10	89.2%	89.9%	92.3%	92.8%	96.4%	96.9%
1.15	90.2%	90.8%	93.2%	93.3%	96.9%	97.2%
1.20	91.3%	91.6%	93.6%	94.1%	97.2%	97.7%
1.25	91.8%	92.6%	93.9%	95.0%	97.8%	98.1%
1.30	92.7%	93.6%	94.2%	95.2%	98.1%	98.5%
1.35	93.4%	93.9%	94.4%	95.4%	98.5%	98.7%
1.40	93.8%	94.2%	94.7%	95.6%	98.6%	98.9%
1.45	94.2%	94.4%	95.2%	95.6%	98.7%	99.5%
1.50	94.5%	95.1%	95.7%	95.7%	98.7%	99.5%
1.55	94.9%	95.3%	95.9%	96.0%	99.0%	99.7%
1.60	95.2%	95.4%	96.2%	96.3%	99.1%	99.7%
1.65	95.6%	95.7%	96.5%	96.3%	99.3%	99.7%
1.70	95.9%	95.8%	96.7%	96.6%	99.5%	99.7%
1.75	96.3%	96.0%	97.0%	96.6%	99.5%	99.7%
1.80	96.4%	96.1%	97.3%	97.1%	99.6%	99.7%
1.85	96.6%	96.4%	97.3%	97.4%	99.6%	99.7%
1.90	96.9%	96.7%	97.3%	97.5%	99.6%	99.7%
1.95	96.9%	96.9%	97.4%	97.8%	99.6%	99.7%
2.00	97.1%	97.2%	97.6%	97.8%	99.6%	99.7%