Estimation of Unknown Parameters Using Partially Observed Data

Eunji Lim

Decision Sciences and Marketing, Adelphi University

Abstract

Purpose: We consider complex stochastic systems, such as supply chains, whose dynamics are controlled by an unknown parameter, such as the arrival or service rates. The purpose of this paper is to provide a simulation-based estimator of the unknown parameter when only partially observed data on the underlying system are available.

Design/methodology/approach: The proposed method treats the unknown parameter as a random variable, and estimates the parameter by computing the conditional expectation of the random variable given the partially observed data. We then express the conditional expectation as a weighted sum of reverse conditional probabilities using Bayes' rule. The reverse conditional probabilities are estimated using simulation.

Findings: Our simulation studies indicate that the proposed estimator converges to the true value of the conditional expectation as the computer time allocated to the simulation increases. The proposed estimator is computed within a few seconds in all of our numerical examples, which demonstrates its time efficiency.

Originality/value: Most of the existing methods for estimating an unknown parameter require a significant amount of simulation, causing long computation delays. The proposed method requires a single simulation run for each candidate of the unknown parameter. Thus, it is designed to carry a significantly reduced computational burden. This feature will enable managers to use the proposed method when making real-time decisions.

Keywords: statistical inference; parameter estimation; simulation; hidden Markov models; supply chain management

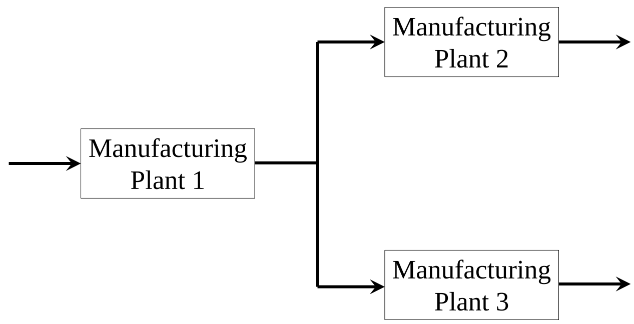

Figure 1: A simplified version of a manufacturer's supply chain

1 Introduction

We consider the problem of estimating an unknown parameter for a stochastic process when only partially observed data on such a process are available. This paper is motivated by a manufacturer whose supply chain consists of multiple manufacturing plants with jobs being transported from one plant to another. Figure 1 depicts a simplified version of such a supply chain. In order to predict the future performance of the supply chain, a computer simulation can be set up and run to provide estimates on future performance. One of the challenges in this setting was that some of the parameters that are necessary to set up such a simulation were not known, and only partial information on the supply chain was collected. For example, the manufacturer needed to know the rate at which jobs were processed at each plant in order to run a computer simulation, but the rate was not known in one of the plants. On the other hand, some partial information, such as the total number of jobs either being processed or waiting in queue at each plant, was observable on a daily basis. Thus, the main problem that the manufacturer faced was how to estimate the unknown parameter based on the partially observed data on his supply chain.

To answer the manufacturer's question, we take the view that the system under consideration can be simulated by updating a vector of real numbers X_t for each $t \in \{\cdots, -1, 0, 1, \cdots\}$, where X_t describes the internal state of the system at time t. Throughout this paper, the stochastic process $X = (X_t : -\infty < t < \infty)$ is assumed to be a Markov chain in the steady state. A stochastic process can be assumed to be in the steady-state, under certain conditions, when it has evolved for a long period of time. It is reasonable to assume that X is in the steady state since, in many applications, the system under consideration has been up and running for a long period of time.

We further assume that only part of X_t is observable and collected as data, and hence, the observable process can be expressed as $Y = (Y_t : -\infty < t < \infty)$, where $Y_t = f(X_t)$ for $-\infty < t < \infty$ and some known function f. For example, X_t can be the total number of jobs at each plant and the amount of time that has passed since service was initiated on a job at each plant, and Y_t can be the numbers of jobs in some (not all) of the plants at time t. In order to simulate the system by updating X_t, we need to know certain parameters such as the rate at which jobs are processed at each plant or the rate at which jobs arrive at each plant. We assume that all of these parameters are known except for one, say θ_*. If we assume that the current time is $t = 0$, the problem can be formulated as an estimation of θ_* based on observed data $Y_0 = y_0, Y_{-1} = y_{-1}, \cdots$.

In our approach, we first treat θ_* as a random variable θ that can take one of the values in $\Theta = \{\theta_1, \cdots, \theta_r\}$. Hence, our best guess on θ_* given $Y_0 = y_0, Y_{-1} = y_{-1}, \cdots$ is the conditional expectation of θ given $Y_0 = y_0, Y_{-1} = y_{-1}, \cdots$, i.e., $\mathbb{E}(\theta | Y_0 = y_0, Y_{-1} = y_{-1}, \cdots)$. We then realize that computing the conditional expectation on the full history $Y_0 = y_0, Y_{-1} = y_{-1}, \cdots$ is computationally challenging and sometimes impractical, and hence, we simplify our problem as an estimation of θ based on $Y_0 = y_0$ only. Our problem thus boils down to how to compute $\mathbb{E}(\theta | Y_0 = y_0)$. We next rewrite $\mathbb{E}(\theta | Y_0 = y_0)$ as follows (using Bayes' rule):

$$\begin{aligned}
\mathbb{E}(\theta | Y_0 = y_0) &= \sum_{i=1}^{r} \theta_i \mathbb{P}(\theta = \theta_i | Y_0 = y_0) \\
&= \sum_{i=1}^{r} \theta_i \left(\frac{\mathbb{P}(Y_0 = y_0 | \theta = \theta_i)\mathbb{P}(\theta = \theta_i)}{\sum_{j=1}^{r} \mathbb{P}(Y_0 = y_0 | \theta = \theta_j)\mathbb{P}(\theta = \theta_j)} \right) \\
&= \frac{\sum_{i=1}^{r} \theta_i \mathbb{P}(Y_0 = y_0 | \theta = \theta_i)\mathbb{P}(\theta = \theta_i)}{\sum_{i=1}^{r} \mathbb{P}(Y_0 = y_0 | \theta = \theta_i)\mathbb{P}(\theta = \theta_i)}. \qquad (1.1)
\end{aligned}$$

We then suggest a way to estimate $\mathbb{P}(\theta = \theta_i)$ and $\mathbb{P}(Y_0 = y_0 | \theta = \theta_i)$ through simulation. By plugging these estimates into (1.1), we will be able to compute an estimator of $\mathbb{E}(\theta | Y_0 = y_0)$.

Typically, a modeler has a rough idea of what the range of possible values of θ is. We therefore assume that Θ is given to the modeler, and he draws a value from Θ uniformly. Assuming that θ is uniformly distributed over Θ, (1.1) is further simplified to

$$\frac{\sum_{i=1}^{r} \theta_i \mathbb{P}(Y_0 = y_0 | \theta = \theta_i)}{\sum_{i=1}^{r} \mathbb{P}(Y_0 = y_0 | \theta = \theta_i)} \qquad (1.2)$$

and it remains to estimate $\mathbb{P}(Y_0 = y_0 | \theta = \theta_i)$ through simulation. There are a couple of challenges in estimating $\mathbb{P}(Y_0 = y_0 | \theta = \theta_i)$ via simulation.

1. How do we sample X_0 from its steady–state distribution? Since $Y_0 = f(X_0)$, one needs to generate X_0 in order to generate Y_0. X_0 is assumed to be in its steady state. In general, the steady–state distribution of X is not available in a closed form. How can we use simulation to sample X_0 from its steady–state distribution?

2. If the event $\{Y_0 = y_0\}$ occurs rarely given $\theta = \theta_i$ for $i \in \{1, \cdots, r\}$, how do we estimate the probability of such an event in a time efficient way?

To answer 1, we use the following procedure. Fix $\theta = \theta_i$. Let $\tilde{X} = (\tilde{X}_t : 1 \le t \le n)$ be a Markov chain with the same transition probabilities as those of X when $\theta_* = \theta_i$. We typically generate $\tilde{X}$ by setting $\tilde{X}_1$ at an arbitrary value and running $\tilde{X}$ forward using the transition probabilities. As n, the length of simulation, gets longer, the distribution of $\tilde{X}_n$ gets closer to the steady–state distribution of X. Once n is large enough so that the distribution of $\tilde{X}_n$ is close to the steady–state distribution of X, one needs to count the number of times $\tilde{Y}_n = f(\tilde{X}_n)$ hits y in order to estimate $\mathbb{P}(Y_0 = y_0|\theta = \theta_i)$. If Y_0 has a large state space and the event $\{Y_0 = y_0\}$ occurs rarely, we may need to generate many replications of $(\tilde{X}_1, \cdots, \tilde{X}_n)$ in order to find an instance where $\tilde{Y}_n = y$. This procedure can be time consuming. To overcome this, we suggest the following remedy. Even in the instance where $\tilde{Y}_n$ does not hit y, we compute the distance between $\tilde{Y}_n$ and y and assign an appropriate weight to such an instance; the closer $\tilde{Y}_n$ is to y, the more weight is assigned to $\tilde{Y}_n$. We express the weight using a function $K : \mathbb{R} \to \mathbb{R}$ called the kernel function. For example, $K(\|w - y\|)$ can be assigned to the instance where $\tilde{Y}_n = w$, where $K(z) = (1/\sqrt{2\pi}) \exp(-z^2/(2\lambda^2))$ for $z \in \mathbb{R}$ and some positive constant λ, or $K(z) = 1/(2h)$ for $-h \le z \le h$ and $K(z) = 0$ otherwise. ($\|(x_1, \cdots, x_d)\| = \sqrt{x_1^2 + \cdots + x_d^2}$ for $(x_1, \cdots, x_d) \in \mathbb{R}^d$.) A natural step to follow is to generate multiple replications of $\tilde{Y}_n$ and compute the average weight assigned to these replications as an estimator of $\mathbb{P}(Y_0 = y_0|\theta = \theta_i)$. However, there is a significant downside to this approach. Simulating $(\tilde{X}_t : 1 \le t \le n)$ multiple times is time consuming. To tackle this issue, we note that generating one more time step in the $\tilde{X}$ process requires much less time than starting a new simulation run of $\tilde{X}$. (For example, it takes an average of 1.67 milliseconds to generate $\tilde{X}_1, \cdots, \tilde{X}_{100}$ in the example in Section 3.3, while it takes an average of 1.68 milliseconds to generate $\tilde{X}_1, \cdots, \tilde{X}_{101}$ in the same example.) Thus, we suggest generating a long single run $(\tilde{X}_t : 1 \le t \le n)$ and use the average of the weights $K(\|\tilde{Y}_1 - y\|), \cdots, K(\|\tilde{Y}_n - y\|)$ as an estimator of $\mathbb{P}(Y_0 = y|\theta = \theta_i)$ instead of generating multiple simulation runs and averaging the weights assigned to $\tilde{Y}_n$ in each run. Plugging the average weight into $\mathbb{P}(Y_0 = y|\theta = \theta_i)$ in (1.2) completes the development of our proposed estimator. Our proposed estimator of $\mathbb{E}(\theta|Y_0 = y)$ is given as follows:

$$\frac{\sum_{i=1}^{r} \theta_i K_i}{\sum_{i=1}^{r} K_i}, \tag{1.3}$$

where $K_i = \sum_{j=1}^{n} K(\|\tilde{Y}_j(\theta_i) - y\|)/n$, $\tilde{Y}_j(\theta_i) = f(X_j(\theta_i))$ for $1 \leq j \leq n$, and $(\tilde{X}_t(\theta_i) : 1 \leq t \leq n)$ is a Markov chain with an arbitrary initial distribution and the same transition probabilities as those of X when $\theta_* = \theta_i$.

Our numerical results in Section 3 illustrate that our proposed estimator shows convergence to $\mathbb{E}(\theta|Y_0 = y_0)$ as n increases. In addition, our proposed estimator is computed within a few seconds in all of our numerical examples, thereby demonstrating its computational efficiency.

In the literature, the problem of estimating an unknown parameter θ_* has been considered in two different settings: (1) the case where a realization of the underlying Markov chain X is observable; and (2) the case where partially observed data Y are observable. In the case where X is observable, the likelihood function method and the Bayesian approach have been used as two main approaches. In the likelihood function method, θ_* is a real parameter that is used in the density function of $(X_{-m}, \cdots, X_0)$ and one should be able to write the likelihood function in terms of the observed data $X_{-m} = x_{-m}, \cdots, X_0 = x_0$ and θ_* only; see Basawa and Prakasa Rao (1980) and the references therein for details. Thus, this approach is not applicable to our setting since we have only partially observed data $Y_{-m} = y_{-m}, \cdots, Y_0 = y_0$ at our disposal and we cannot express the likelihood function using $y_{-m}, \cdots, y_0$ only. On the other hand, in the Bayesian approach, θ_* is viewed as a random variable θ that takes one of the values in Θ. After observing $X_{-m} = x_{-m}, \cdots, X_0 = x_0$, one can compute the conditional expectation of θ given $X_{-m} = x_{-m}, \cdots, X_0 = x_0$ using Bayes' rule as follows:

$$\mathbb{E}(\theta|X_{-m} = x_{-m}, \cdots, X_0 = x_0) = \frac{\sum_{i=1}^{r} \theta_i \mathbb{P}(X_{-m} = x_{-m}, \cdots, X_0 = x_0|\theta = \theta_i)\mathbb{P}(\theta = \theta_i)}{\sum_{i=1}^{r} \mathbb{P}(X_{-m} = x_{-m}, \cdots, X_0 = x_0|\theta = \theta_i)\mathbb{P}(\theta = \theta_i)}.$$

Therefore, our proposed approach can be viewed as an extension of Bayes' approach to the case where partially observed data are available. Furthermore, we recognize that the computational burden of computing $\mathbb{P}(Y_{-m} = y_{-m}, \cdots, Y_0 = y_0|\theta = \theta_i)$ is heavy, we suggest a fast and efficient algorithm for estimating $\mathbb{P}(Y_0 = y_0|\theta = \theta_i)$ through simulation.

In the case where Y is observable, many researchers have focused on the filtering approach, which is often referred to as the sequential Monte Carlo method, in order to estimate $\mathbb{E}(\theta|Y_{-m} = y_{-m}, \cdots, Y_0 = y_0)$. The filtering approach was originally proposed to compute the conditional expectation of a certain random variable (rather than a parameter) given

$Y_{-m} = y_{-m}, \cdots, Y_0 = y_0$; see Liu and Chen (1995), Casella and Robert (1996), Kitagawa (1996), Beadle and Djurić (1997), Liu and Chen (1998), Pitt and Shephard (1999), Carpenter et al. (1999), Pitt and Shephard (1999), Liu (2001), Gilks and Berzuini (2001), Fox (2003), Künsch (2005), Doucet et al. (2006), Doucet and Johansen (2011), Bhada and Ionides (2014), Leippold and Yang (2019), and Jacob et al. (2019) for various filtering-based methods. Recently, the filtering approach has been modified so that one can estimate a parameter given partially observed data (Kantas et al. (2009), Nemeth et al. (2014), Yang et al. (2018)). Even though filtering-based methods have well-established theories in the literature, one of their drawbacks is their heavy computational burdens. The filtering approach proceeds by generating l copies of X_t (for each $t \in \{-m, \cdots, 0\}$), and each copy requires running an independent simulation. To ensure the convergence of a filtering-based estimator, l needs to be large, thereby placing a heavy computational burden on the modeler's side. On the other hand, the proposed method requires a single simulation run for each candidate of the unknown parameter. Thus, it carries a significantly reduced computational burden. This feature will enable managers to use the proposed method when making real-time decisions.

This paper is organized as follows. The proposed method is precisely described in Section 2. In Section 3, we apply the proposed method in two different settings. In Section 3.1, we consider a Jackson network with three stations, where $\mathbb{E}(\theta|Y_0 = y_0)$ is available in a closed form. Since $\mathbb{E}(\theta|Y_0 = y_0)$ is available in a closed form, we can compare our proposed estimator to the true value of $\mathbb{E}(\theta|Y_0 = y_0)$. The results in Section 3.1 indicate that our proposed estimator successfully converges to the true value of $\mathbb{E}(\theta|Y_0 = y_0)$ as n increases. In Section 3.2, we consider a more realistic example, where $\mathbb{E}(\theta|Y_0 = y_0)$ is not available in a closed form. The results in Section 3.2 illustrate that our proposed estimator shows convergence as n increases. The proposed estimators are computed within a few seconds in all of our numerical examples. While conducting numerical experiments, we observed that $\mathbb{E}(\theta|Y_0 = y_0, Y_{-1} = y_{-1}, \cdots)$ often coincides with the true value of θ_*. In Section 4, we investigate this issue further and prove that $\mathbb{E}(\theta|Y_0 = y_0, Y_{-1} = y_{-1}, \cdots, Y_{-m} = y_{-m}) \to \theta_*$ as $m \to \infty$ in a specific setting where we can observe the total number of jobs in a network of queues with an unknown mean interarrival time. Concluding remarks are included in Section 5.

2 Proposed Method

In this section, we describe the proposed method more precisely. We assume that the system under consideration can be simulated by recursively updating some state variables, so it can be modeled as a general state space Markov chain $X = (X_t : -\infty < t < \infty)$ in the steady state and the observable process is $Y = (Y_t : -\infty < t < \infty)$, where $Y_t = f(X_t)$ for some known function f and $t \in \{\cdots, -1, 0, 1, \cdots\}$. The transition kernel of X depends on some parameters, one of which is unknown to the modeler. We call the unknown parameter $\theta_* \in \mathbb{R}$. We further assume that the current time is $t = 0$, and some partially observed data $Y_0 = y_0, Y_{-1} = y_{-1}, \cdots$ are available. In this setting, our goal is to estimate the unknown parameter θ_* using observed data. We take the view that a modeler has an idea of what the range of possible values of θ_* is. Thus, the set $\Theta = \{\theta_1, \cdots, \theta_r\}$ of possible values of θ_* is assumed to be given to the modeler. We treat θ_* as a random variable θ (rather than a static parameter) that takes one of the values in Θ. To simplify the problem, we consider conditioning only on the most recent observation $Y_0 = y_0$. Then, our best guess on θ_* is the conditional expectation of θ given $Y_0 = y_0$, i.e., $\mathbb{E}(\theta|Y_0 = y_0)$.

The question boils down to how to estimate $\mathbb{E}(\theta|Y_0 = y_0)$. Using Bayes' rule, we rewrite $\mathbb{E}(\theta|Y_0 = y_0)$ as

$$\mathbb{E}(\theta|Y_0 = y_0) = \frac{\sum_{i=1}^r \theta_i \mathbb{P}(Y_0 = y_0|\theta = \theta_i)\mathbb{P}(\theta = \theta_i)}{\sum_{i=1}^r \mathbb{P}(Y_0 = y_0|\theta = \theta_i)\mathbb{P}(\theta = \theta_i)},$$

which is further simplified to

$$\mathbb{E}(\theta|Y_0 = y_0) = \frac{\sum_{i=1}^r \theta_i \mathbb{P}(Y_0 = y_0|\theta = \theta_i)}{\sum_{i=1}^r \mathbb{P}(Y_0 = y_0|\theta = \theta_i)}, \tag{2.1}$$

if the modeler draws θ from Θ uniformly.

The key feature of the proposed method is that we estimate $\mathbb{E}(Y_0 = y_0|\theta = \theta_i)$ using

$$K_i = \frac{1}{n}\sum_{j=1}^n K(\|\tilde{Y}_j(\theta_i) - y\|),$$

where $\tilde{Y}_j(\theta_i) = f(X_j(\theta_i))$ for $1 \leq j \leq n$, and $(\tilde{X}_t(\theta_i) : 1 \leq t \leq n)$ is a Markov chain with an arbitrary initial distribution and the same transition kernel as that of X when $\theta_* = \theta_i$.

By replacing $\mathbb{E}(Y_0 = y_0|\theta = \theta_i)$ in (2.1) with K_i, we obtain the proposed estimator $\hat{\theta}_n$ as follows:

$$\hat{\theta}_n = \frac{\sum_{i=1}^r \theta_i K_i}{\sum_{i=1}^r K_i}$$

if the modeler chooses $\theta = \theta_i$ uniformly from Θ, or

$$\hat{\theta}_n = \frac{\sum_{i=1}^r \theta_i K_i p(\theta_i)}{\sum_{i=1}^r K_i p(\theta_i)}$$

if the modeler chooses $\theta = \theta_i$ with probability $p(\theta_i)$ for $1 \leq i \leq r$.

The proposed estimator can be computed using the following algorithm.

Proposed Method

Step 0. Let $\Theta = \{\theta_1, \cdots, \theta_r\}$ and the kernel function $K : \mathbb{R} \to \mathbb{R}$ be given.

Step 1. Set $i = 1$.

Step 2. Set $\theta = \theta_i$. Take a value x_0 randomly from the state space of X and set $\tilde{X}_0(\theta_i) = x_0$.

Step 3. Generate $\tilde{X}_1(\theta_i), \cdots, \tilde{X}_n(\theta_i)$ using the same transition probabilities as those of X when $\theta = \theta_i$. Compute $\tilde{Y}_1(\theta_i), \cdots, \tilde{Y}_n(\theta_i)$ using $\tilde{Y}_j(\theta_i) = f(\tilde{X}_j(\theta_i))$ for $1 \leq j \leq n$.

Step 4. Compute

$$K_i = \frac{1}{n} \sum_{j=1}^n K(\|\tilde{Y}_j(\theta_i) - y\|),$$

where $\|(x_1, \cdots, x_d)\| = \sqrt{x_1^1 + \cdots + x_d^2}$ for $(x_1, \cdots, x_d) \in \mathbb{R}^d$.

Step 5. If $i < r$, increase i by 1 and repeat Steps 2 through 4. If $i = r$, compute the proposed estimator $\hat{\theta}_n$ as follows:

$$\hat{\theta}_n = \frac{\sum_{i=1}^r \theta_i K_i}{\sum_{i=1}^r K_i}.$$

3 Numerical Results

In this section, we apply the proposed method in two different settings. In Section 3.1, we consider a Jackson network (a network of queues with exponentially distributed service times and interarrival times) with three stations, where $\mathbb{E}(\theta | Y_0 = y_0)$ is available in a closed form.

Since $\mathbb{E}(\theta|Y_0 = y_0)$ is available in a closed form, we can compare the proposed estimator to the true value of $\mathbb{E}(\theta|Y_0 = y_0)$. In Section 3.2, we turn to a more realistic example, where we consider a network with three stations whose service times are not exponentially distributed, so $\mathbb{E}(\theta|Y_0 = y_0)$ is not available in a closed form.

3.1 Jackson Network with an Unknown Arrival Rate

We consider a network with three stations, say Stations 1, 2, and 3. Jobs arrive externally with the interarrival times exponentially distributed with a mean of θ_*. Once a job arrives externally, it joins Station 1. Each station has a single server and infinite buffer capacity. The service times at Stations 1, 2, and 3 are independent of each other and follow exponential distributions with means of 4, 6, and 8.5, respectively. The service times are independent of the interarrival times. A job leaving Station 1 goes to Station 2 or 3 with probabilities 0.6 and 0.4, respectively. A job leaving Station 2 or 3 leaves the network forever. We let $X_t = (N_t(1), N_t(2), N_t(3))$, where $N_t(1), N_t(2)$ and $N_t(3)$ are the numbers of jobs at Stations 1, 2, and 3 at the tth departure or arrival epoch of jobs. It should be noted that $X = (X_t : -\infty < t < \infty)$ is a general state space Markov chain.

We consider the situation where we can observe the numbers of jobs in Stations 1 and 2 only, and the number of jobs in Station 3 is not observable. Hence, the observable process is $Y = (Y_t : -\infty < t < \infty) = ((N_t(1), N_t(2)) : -\infty < t < \infty)$ and $Y_t = f(X_t)$, where $f : \mathbb{R}^3 \to \mathbb{R}^2$ is defined by $f(x_1, x_2, x_3) = (x_1, x_2)$ for $(x_1, x_2, x_3) \in \mathbb{R}^3$.

Assuming that the current time is $t = 0$, our goal is to estimate $\mathbb{E}(\theta|Y_0 = (16, 4))$.

We assume that the set of possible values for θ_* is given by $\Theta = \{4.1, 4.3, 4.5, 4.7\}$. To compute the proposed estimator $\hat{\theta}_n$, we applied the proposed method in Section 2 with $x_0 = (0, 0, 0)$ and $K = K_n$, which is defined by

$$K_n(z) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{z^2}{2\lambda_n^2})$$

for $z \in \mathbb{R}$, where $\lambda_n = 3/\lceil n/1000 \rceil$ and $\lceil x \rceil$ is the smallest integer greater than or equal to $x \in \mathbb{R}$.

We repeated this procedure 100 times and obtained 100 independent and identically distributed (iid) copies of our estimator $\hat{\theta}_n$. Based on these 100 copies, we computed the 95% confidence interval of $\hat{\theta}_n$ and the average time required to generate a copy of $\hat{\theta}_n$. Table 1 reports these 95% confidence intervals and the average times for a variety of n values. The true value of $\mathbb{E}(\theta|Y_0 = (16,4))$ can be computed using (1.2) and the fact that

$$\mathbb{P}(Y_0 = (16,4)|\theta = \theta_i) = (1-\theta_i/4)(\theta_i/4)^{16}(1-\theta_i/6)(\theta_i/4)^4$$

for $1 \le i \le 4$; see, for example, Equation (2.5) on page 17 of Chen and Yao (2001). The true value of $\mathbb{E}(\theta|Y_0 = (16,4))$ is reported in the last row of Table 1. Table 1 shows that the proposed estimator $\hat{\theta}_n$ converges to the true value of $\mathbb{E}(\theta|Y_0 = (16,4)))$ as n increases.

[INSERT TABLE 1 HERE.]

3.2 Jackson Network with an Unknown Service Rate

We consider the same problem setting as in Section 3.1, except that we now assume that

- the mean interarrival time is known to be 4.3,
- the mean service time at Station 1 is the unknown parameter θ_*. The set of possible values of θ_* is given by $\Theta = \{3.6, 3.8, 4.0, 4.2\}$, and
- we can only observe the number of jobs in Station 1, i.e., the observable process is $Y = (Y_t : -\infty < t < \infty) = (N_t(1) : -\infty < t < \infty)$ and $Y_t = f(X_t)$, where $f : \mathbb{R}^3 \to \mathbb{R}$ is given by $f(x_1, x_2, x_3) = x_1$ for $(x_1, x_2, x_3) \in \mathbb{R}^3$.

We generated 100 iid copies of our estimator $\hat{\theta}_n$ of $\mathbb{E}(\theta|Y_0 = 16)$. Based on these 100 copies, we computed the 95% confidence intervals of $\hat{\theta}_n$ and the average time required to generate a copy of $\hat{\theta}_n$. Table 2 reports these 95% confidence intervals and the average times for a variety of n values. The true value of $\mathbb{E}(\theta|Y_0 = 16)$ is computed in a similar way to Section 6.1 and is reported in the last row of Table 2. Table 2 shows that the proposed estimator $\hat{\theta}_n$ converges to the true value of $\mathbb{E}(\theta|Y_0 = 16)$ as n increases.

[INSERT TABLE 2 HERE.]

3.3 Non-Jackson Network with an Unknown Arrival Rate

We consider a network with three stations, say Stations 1, 2, and 3. Jobs arrive externally with the interarrival times uniformly distributed over $\{\theta_* - 1, \theta_*, \theta_* + 1\}$. Thus, the mean interarrival time is the unknown parameter θ_*. Once a job arrives externally, it joins Station 1. Each station has a single server with a buffer capacity of 40. The service times at Stations 1, 2, and 3 are independent of each other and are uniformly distributed over $\{4, 5, 6\}$, $\{6, 7, 8, 9, 10\}$, and $\{10, 11, 12, 13, 14, 15\}$, respectively. The service times are independent of the interarrival times. A job leaving Station 1 goes to Station 2 or 3 with probabilities 0.6 and 0.4, respectively. A job leaving Station 2 or 3 leaves the network forever. We let $X_t = (N_t(1), N_t(2), N_t(3), S_t(1), S_t(2), S_t(3), A_t)$, where $N_t(1), N_t(2)$, and $N_t(3)$ are the numbers of jobs at Stations 1, 2, and 3 at the beginning of the tth time period, respectively, and $S_t(1), S_t(2)$, and $S_t(3)$ are the amounts of time that have passed since service was initiated in Stations 1, 2, and 3, respectively (if there is no job at the server of Station i for $i = 1, 2, 3$, $S_t(i)$ is set to be -1), and A_t is the amount of time that has passed since the last job arrived externally. It should be noted that $X = (X_t : -\infty < t < \infty)$ is a general state space Markov chain.

We consider the situation where we can only observe the numbers of jobs in Stations 1, 2, and 3. The elapsed service times and the elapsed interarrival time are not observable. Hence, the observable process is $Y = (Y_t : -\infty < t < \infty) = ((N_t(1), N_t(2), N_t(3)) : -\infty < t < \infty)$ and $Y_t = f(X_t)$, where $f : \mathbb{R}^7 \to \mathbb{R}^3$ is given by $f(x_1, x_2, x_3, x_4, x_5, x_6, x_7) = (x_1, x_2, x_3)$ for $(x_1, x_2, x_3, x_4, x_5, x_6, x_7) \in \mathbb{R}^7$.

Assuming that the current time is $t = 0$, our goal is to estimate $\mathbb{E}(\theta|Y_0 = (12, 6, 13))$. We assume that the set of possible values for θ_* is given by $\Theta = \{4, 5, 6, 7, 8\}$. To compute the proposed estimator $\hat{\theta}_n$, we applied the proposed method in Section 2 in the same manner as in Section 3.1. We generated 100 iid copies of $\hat{\theta}_n$. Based on these 100 copies, we computed the 95% confidence interval of $\hat{\theta}_n$ and the average time required to generate a copy of $\hat{\theta}_n$. Table 3 reports these 95% confidence intervals and the average times for a variety of n values. In Table 4, we set $\Theta = \{2, 3, 4, 5, 6, 7, 8\}$ and repeated the above procedure to obtain the proposed estimator $\hat{\theta}_n$ of $\mathbb{E}(\theta|Y_0 = (18, 38, 10))$ for a variety of n values.

In both Tables 3 and 4, the proposed estimator is computed in a few seconds and shows convergence as n increases.

[INSERT TABLES 3 AND 4 HERE.]

3.4 Non-Jackson Network with an Unknown Service Rate

We consider the same problem setting as in Section 3.3, except that we now assume that

- the interarrival times are known to follow a uniform distribution over $\{4, 5, 6\}$,
- the mean service time at Station 1 is the unknown parameter θ_*. The service times at Station 1 follow a uniform distribution over $\{\theta_* - 1, \theta_*, \theta_* + 1\}$. The set of possible values of θ_* is given by $\Theta = \{4, 5, 6, 7, 8\}$, and

- we can only observe the number of jobs at Station 1, i.e., $Y_t = N_t(1)$ for $-\infty < t < \infty$ and $Y_t = f(X_t)$, where $f : \mathbb{R}^7 \to \mathbb{R}$ is given by $f(x_1, x_2, x_3, x_4, x_5, x_6, x_7) = x_1$ for $(x_1, x_2, x_3, x_4, x_5, x_6, x_7) \in \mathbb{R}^7$.

We generated 100 iid copies of the proposed estimator $\hat{\theta}_n$ of $\mathbb{E}(\theta|Y_0 = 12)$. Based on these 100 copies, we computed the 95% confidence interval of $\hat{\theta}_n$ and the average time required to generate a copy of $\hat{\theta}_n$. Table 5 reports these 95% confidence intervals and the average times for a variety of n values. In Table 6, we set $\Theta = \{2, 3, 4, 5, 6, 7, 8\}$ and repeated the above procedure to obtain the proposed estimator $\hat{\theta}_n$ of $\mathbb{E}(\theta|Y_0 = 18)$ for a variety of n values.

In both Tables 5 and 6, the proposed estimator is computed in a few seconds and shows convergence as n increases.

[INSERT TABLES 5 AND 6 HERE.]

4 Justification of the Conditional Expectation

While conducting numerical experiments, we observed that $\mathbb{E}(\theta|Y_0 = y_0, Y_{-1} = y_{-1}, \cdots)$ often coincides with the true value of θ_*. In this section, we investigate this issue further and prove that $\mathbb{E}(\theta|Y_0 = y_0, Y_{-1} = y_{-1}, \cdots, Y_{-m} = y_{-m}) \to \theta_*$ as $m \to \infty$ in a specific setting. More precisely, we consider a situation where we observe the total number of jobs in a network of queues with an unknown interarrival rate. The system under consideration can be described as follows.

We consider a network with J stations, say Stations 1 through J, each with one server and N_j waiting rooms for $1 \le j \le J$. The service times of jobs at Station j are iid, following an exponential distribution with a mean of $1/\lambda_j$ for $j = 1, \cdots, J$. Jobs travel among the stations following a routing matrix $(p_{jk} : j, k = 1, \cdots, J)$, where p_{jk} is the probability that a job leaving Station j will go to Station k for $1 \le j, k \le J$. At each station, all jobs are served on a first-come-first-served basis. If a job entering a station finds that the server and all the waiting rooms are occupied, then it leaves the network forever. Jobs arrive from outside, following a Poisson process with rate $\theta_* > 0$. Each arrival is independently routed to Station j with probability $p_{0j} > 0$ and $p_{01} + \cdots + p_{0J} = 1$. For $t \in \{\cdots, -1, 0, 1, \cdots\}$ and $j = \{1, \cdots, J\}$, let $X_t(j)$ be the number of jobs at Station j at the tth departure or arrival epoch of the jobs. Then $X = (X_t : -\infty < t < \infty) = ((X_t(1), \cdots, X_t(J)) : -\infty < t < \infty)$ is a Markov chain with a finite state space. We assume that we can only observe the total number of jobs in the network, so Y_t is the total number of jobs in the network at the tth transition epoch. We also assume that $\lambda_1, \cdots, \lambda_J$, $(p_{jk} : 1 \leq j, k \leq J)$, and p_{0j} for $1 \leq j \leq J$ are known, but θ_* is unknown.

The following assumption is needed.

A1: Assume that each station may both receive external input and deliver external output (possibly via other stations), i.e., for each j (i) either $p_{0j} > 0$ or $p_{0l_1} p_{l_1 l_2} \cdots p_{l_v j}$ for some $l_1, \cdots, l_v$ and (ii) either $1 - (p_{j1} + \cdots + p_{jJ}) > 0$ or $p_{jl_1} p_{l_1 l_2} \cdots (1 - (p_{l_v 1} + \cdots + p_{l_v J})) > 0$ for some $l_1, \cdots, l_v$.

A1 implies the irreducibility of $(X_t : -\infty < t < \infty)$. Since X is a finite-state Markov chain, it is positive recurrent. Hence, the irreducibility of X implies the existence of a unique stationary measure, say π_{θ_*}. We need the following two additional assumptions.

A2: X is in the steady state, i.e., X_t is distributed according to its stationary distribution π_{θ_*} for $t \in \{\cdots, -1, 0, 1, \cdots\}$.

A3: θ is a random variable that takes on one of the values in $\Theta = \{\theta_1, \cdots, \theta_r\}$, where one of $\theta_1, \cdots, \theta_r$ is equal to θ_*.

Under A1, A2, and A3, the following theorems establish $\mathbb{E}(\theta | Y_0 = y_0, \cdots, Y_{-m} = y_{-m}) \to \theta_*$ as $m \to \infty$, justifying the use of the conditional expectation as an estimator of θ_*. In the proof of Theorem 1, $\mathbb{E}_\theta$ and $\mathbb{P}_\theta$ denote the conditional expectation given $\theta_* = \theta$ and the conditional probability given $\theta_* = \theta$, respectively.

Theorem 1. *Assume A1–A3. If $\theta_k \neq \theta_*$ for $k = 1, \cdots, r$,*

$$\mathbb{P}(\theta = \theta_k | Y_0 = y_0, \cdots, Y_{-m} = y_{-m}) \to 0$$

as $m \to \infty$. Also,

$$\mathbb{P}(\theta = \theta_* | Y_0 = y_0, \cdots, Y_{-m} = y_{-m}) \to 1$$

as $m \to \infty$.

Proof of Theorem 1. Since

$$\mathbb{P}(\theta = \theta_k | Y_0 = y_0, \cdots, Y_{-m} = y_{-m}) = \frac{\mathbb{P}_{\theta_k}(Y_0 = y_0, \cdots, Y_{-m} = y_{-m}) \mathbb{P}(\theta = \theta_k)}{\sum_{j=1}^r \mathbb{P}_{\theta_j}(Y_0 = y_0, \cdots, Y_{-m} = y_{-m}) \mathbb{P}(\theta = \theta_j)},$$

it suffices to show that

$$\frac{\mathbb{P}_{\theta_k}(Y_0 = y_0, \cdots, Y_{-m} = y_{-m})}{\mathbb{P}_{\theta_*}(Y_0 = y_0, \cdots, Y_{-m} = y_{-m})} \to 0 \tag{4.1}$$

as $m \to \infty$ for $\theta_k \neq \theta_*$.

To simplify the notation, we will focus on the case where there are only 2 stations in the network, i.e., $J = 2$. Generalization to the case where $J > 2$ is straightforward.

For any nonnegative integers j and k, we denote the stationary probability that $X_t(1) = j$ and $X_t(2) = k$ given that the arrival rate equals θ by $\pi_\theta(j, k)$.

We first note that $(Y_t : -\infty < t < \infty)$ itself is a time-homogeneous Markov chain with the following transition probabilities:

$$\begin{aligned}
\mathbb{P}_\theta(Y_{t+1} = 1|Y_t = 0) &= 1 \\
\mathbb{P}_\theta(Y_{t+1} = 0|Y_t = 1) &= \pi_\theta(1,0)\frac{\lambda_1(1 - p_{11} - p_{12})}{\theta + \lambda_1} + \pi_\theta(0,1)\frac{\lambda_2(1 - p_{21} - p_{22})}{\theta + \lambda_2} \\
\mathbb{P}_\theta(Y_{t+1} = 1|Y_t = 1) &= \pi_\theta(1,0)\frac{\lambda_1(p_{11} + p_{12})}{\theta + \lambda_1} + \pi_\theta(0,1)\frac{\lambda_2(p_{21} + p_{22})}{\theta + \lambda_2} \\
\mathbb{P}_\theta(Y_{t+1} = 2|Y_t = 1) &= \pi_\theta(1,0)\frac{\theta}{\theta + \lambda_1} + \pi_\theta(0,1)\frac{\theta}{\theta + \lambda_2}.
\end{aligned}$$

For $j \geq 2$,

$$\begin{aligned}
\mathbb{P}_\theta(Y_{t+1} = j-1|Y_t = j) &= \sum_{k=1}^{N_1+1} \pi_\theta(k,0)\frac{\lambda_1(1 - p_{11} - p_{12})}{\theta + \lambda_1} + \sum_{k=1}^{N_2+1} \pi_\theta(0,k)\frac{\lambda_2(1 - p_{21} - p_{22})}{\theta + \lambda_2} \\
&\quad + \left(\sum_{l=1}^{N_2+1}\sum_{k=1}^{N_1+1} \pi_\theta(k,l)\right)\frac{\lambda_1(1 - p_{11} - p_{12}) + \lambda_2(1 - p_{21} - p_{22})}{\theta + \lambda_1 + \lambda_2} \\
\mathbb{P}_\theta(Y_{t+1} = j|Y_t = j) &= \sum_{k=1}^{N_1+1} \pi_\theta(k,0)\frac{\lambda_1(p_{11} + p_{12})}{\theta + \lambda_1} + \sum_{k=1}^{N_2+1} \pi_\theta(0,k)\frac{\lambda_2(p_{21} + p_{22})}{\theta + \lambda_2} \\
&\quad + \left(\sum_{l=1}^{N_2+1}\sum_{k=1}^{N_1+1} \pi_\theta(k,l)\right)\frac{\lambda_1(p_{11} + p_{12}) + \lambda_2(p_{21} + p_{22})}{\theta + \lambda_1 + \lambda_2} \\
\mathbb{P}_\theta(Y_{t+1} = j+1|Y_t = j) &= \sum_{k=1}^{N_1+1} \pi_\theta(k,0)\frac{\theta}{\theta + \lambda_1} + \sum_{k=1}^{N_2+1} \pi_\theta(0,k)\frac{\theta}{\theta + \lambda_2} \\
&\quad + \left(\sum_{l=1}^{N_2+1}\sum_{k=1}^{N_1+1} \pi_\theta(k,l)\right)\frac{\theta}{\theta + \lambda_1 + \lambda_2}.
\end{aligned}$$

For $j \geq 0$, we will denote

$$\begin{aligned}
\mathbb{P}_\theta(Y_{t+1} = j-1|Y_t = j) &\triangleq P_\theta(j,-) \\
\mathbb{P}_\theta(Y_{t+1} = j|Y_t = j) &\triangleq P_\theta(j,0) \\
\mathbb{P}_\theta(Y_{t+1} = j+1|Y_t = j) &\triangleq P_\theta(j,+).
\end{aligned}$$

and denote the stationary distribution of $(Y_t : -\infty < t < \infty)$ by π_{y,θ_*} when the arrival rate is θ_*.

Now we note that by Theorem 1.23 on page 26 of Durrett (2011), for each $j \geq 0$, we have

$$\begin{aligned} N_{j+}/s &\to \pi_{y,\theta_*}(j)P_{\theta_*}(j,+) \\ N_{j0}/s &\to \pi_{y,\theta_*}(j)P_{\theta_*}(j,0) \\ N_{j-}/s &\to \pi_{y,\theta_*}(j)P_{\theta_*}(j,-) \end{aligned} \tag{4.2}$$

almost surely as $s \to \infty$, where given that the arrival rate is θ_*, N_{j+} is the number of times $Y_k = j$ and $Y_{k+1} = Y_k + 1$ for $-m \leq k \leq -1$, N_{j0} is the number of times $Y_k = j$ and $Y_{k+1} = Y_k$ for $-m \leq k \leq -1$, and N_{j-} is the number of times $Y_k = j$ and $Y_{k+1} = Y_k - 1$ for $-m \leq k \leq -1$, i.e.,

$$\begin{aligned} N_{j+} &= \sum_{k=-m}^{-1} I(Y_k = j, Y_{k+1} = Y_k + 1 | \theta = \theta_*), \\ N_{j0} &= \sum_{k=-m}^{-1} I(Y_k = j, Y_{k+1} = Y_k | \theta = \theta_*), \text{ and} \\ N_{j-} &= \sum_{k=-m}^{-1} I(Y_k = j, Y_{k+1} = Y_k - 1 | \theta = \theta_*). \end{aligned}$$

Without loss of generality, we can assume that we have a sample path $y_{-m}, \cdots, y_0$ where (4.2) is satisfied.

Next, we note that

$$\mathbb{P}_\theta(Y_0 = y_0, \cdots, Y_{-m} = y_{-m} | \theta) = \mathbb{P}_\theta(Y_0 = y_0) \prod_{j=0}^{N_1+N_2+2} P_\theta(j,+)^{N_{j+}} P_\theta(j,0)^{N_{j0}} P_\theta(j,-)^{N_{j-}}$$

and that

$$\begin{aligned} \frac{\mathbb{P}_\theta(Y_0 = y_0, \cdots, Y_{-m} = y_{-m})}{\mathbb{P}_{\theta_*}(Y_0 = y_0, \cdots, Y_{-m} = y_{-m})} &= \frac{\mathbb{P}_\theta(Y_0 = y_0)}{\mathbb{P}_{\theta_*}(Y_0 = y_0)} \\ &\cdot \prod_{j=0}^{N_1+N_2+2} \left(\frac{P_\theta(j,+)}{P_{\theta_*}(j,+)}\right)^{N_{j+}} \left(\frac{P_\theta(j,0)}{P_{\theta_*}(j,0)}\right)^{N_{j0}} \left(\frac{P_\theta(j,-)}{P_{\theta_*}(j,-)}\right)^{N_{j-}}. \end{aligned}$$

To prove (4.1), it suffices to show that for each $j \geq 1$ and $\theta \neq \theta_*$,

$$(1/m) \log \left(\frac{P_\theta(j,+)}{P_{\theta_*}(j,+)}\right)^{N_{j+}} \left(\frac{P_\theta(j,0)}{P_{\theta_*}(j,0)}\right)^{N_{j0}} \left(\frac{P_\theta(j,-)}{P_{\theta_*}(j,-)}\right)^{N_{j-}} \to a_j \tag{4.3}$$

as $m \to \infty$ for some negative number a_j.

To prove (4.3), we note that

$$\begin{aligned}
&\frac{1}{m}\log\left(\frac{P_\theta(j,+)}{P_{\theta_*}(j,+)}\right)^{N_{j+}}\left(\frac{P_\theta(j,0)}{P_{\theta_*}(j,0)}\right)^{N_{j0}}\left(\frac{P_\theta(j,-)}{P_{\theta_*}(j,-)}\right)^{N_{j-}} \\
&= \frac{N_{j+}}{m}\log\left(\frac{P_\theta(j,+)}{P_{\theta_*}(j,+)}\right) + \frac{N_{j0}}{m}\log\left(\frac{P_\theta(j,0)}{P_{\theta_*}(j,0)}\right) + \frac{N_{j-}}{m}\log\left(\frac{P_\theta(j,-)}{P_{\theta_*}(j,-)}\right) \\
&\to \pi_{y,\theta_*}(j)P_{\theta_*}(j,+)\log\left(\frac{P_\theta(j,+)}{P_{\theta_*}(j,+)}\right) + \pi_{y,\theta_*}(j)P_{\theta_*}(j,0)\log\left(\frac{P_\theta(j,0)}{P_{\theta_*}(j,0)}\right) \\
&\qquad + \pi_{y,\theta_*}(j)P_{\theta_*}(j,-)\log\left(\frac{P_\theta(j,-)}{P_{\theta_*}(j,-)}\right) \\
&= \pi_{y,\theta_*}(j)\log\left(\frac{P_\theta(j,+)}{P_{\theta_*}(j,+)}\right)^{P_{\theta_*}(j,+)}\left(\frac{P_\theta(j,0)}{P_{\theta_*}(j,0)}\right)^{P_{\theta_*}(j,0)} \\
&\qquad \cdot\left(\frac{1-P_\theta(j,+)-P_\theta(j,0)}{1-P_{\theta_*}(j,+)-P_{\theta_*}(j,0)}\right)^{1-P_{\theta_*}(j,+)-P_{\theta_*}(j,0)} \\
&\triangleq \rho(\theta)
\end{aligned}$$

as $m \to \infty$ almost surely.

We next prove that $\rho(\theta)$ is increasing when $\theta < \theta_*$, decreasing when $\theta > \theta_*$, and zero when $\theta = \theta_*$. To prove this, we note that

$$\begin{aligned}
\frac{d\rho}{d\theta} &= \pi_{y,\theta_*}(j)\left(\frac{P_{\theta_*}(j,+)}{P_\theta(j,+)}\frac{dP_\theta(j,+)}{d\theta} + \frac{P_{\theta_*}(j,0)}{P_\theta(j,0)}\frac{dP_\theta(j,0)}{d\theta}\right. \\
&\qquad \left. -\left(\frac{1-P_{\theta_*}(j,+)-P_{\theta_*}(j,0)}{1-P_\theta(j,+)-P_\theta(j,0)}\right)\left(\frac{dP_\theta(j,+)}{d\theta}+\frac{dP_\theta(j,0)}{d\theta}\right)\right),
\end{aligned}$$

$P_{\theta_*}(j,+)$ is a decreasing function of θ, $P_{\theta_*}(j,0)$ is a decreasing function of θ, $1 - P_{\theta_*}(j,+) - P_{\theta_*}(j,0)$ is an increasing function of θ, and hence,

$$\begin{aligned}
&\frac{d\rho}{d\theta} > 0 \text{ when } \theta < \theta_* \\
&\frac{d\rho}{d\theta} < 0 \text{ when } \theta_* < \theta, \text{ and} \\
&\frac{d\rho}{d\theta} = 0 \text{ when } \theta = \theta_*.
\end{aligned}$$

Since $\rho(\theta) \downarrow -\infty$ as $\theta \downarrow 0$, $\rho(\theta) \downarrow -\infty$ as $\theta \uparrow \infty$, we can conclude that $\rho(\theta)$ is increasing when $\theta < \theta_*$, decreasing when $\theta > \theta_*$, and zero when $\theta = \theta_*$, and hence, (4.3) follows. □

Theorem 2. *Under A1, A2, and A3,*

$$\mathbb{E}(\theta|Y_0 = y_0, \cdots, Y_{-m} = y_{-m}) \to \theta_*$$

as $m \to \infty$ a.s.

Proof of Theorem 2. By Theorem 1,

$$\mathbb{E}(\theta|Y_0 = y_0, \cdots, Y_{-m} = y_{-m}) = \sum_{i=1}^{r} \theta_i \mathbb{P}(\theta = \theta_i|Y_0 = y_0, \cdots, Y_{-m} = y_{-m}) \to \theta_*$$

as $m \to \infty$, which completes the proof of Theorem 2. □

Remark 1. *Theorem 2 states that the conditional expectation $\mathbb{E}(\theta|Y_0 = y_0, \cdots, Y_{-m} = y_{-m})$ converges to the true value of the unknown parameter θ_* as the length of the historical data increases to infinity. This, in turn, implies that the proposed algorithm, when modified to compute the conditional expectation given the full history of Y, computes the true value of the unknown parameter.*

5 Conclusions

In this paper, we proposed an efficient algorithm for estimating an unknown parameter θ_* that is required to describe the dynamics of a stochastic process $X = (X_t : -\infty < t < \infty)$ when partially observed data $Y_0 = y_0, Y_{-1} = y_{-1}, \cdots$ are available. Our proposed method treats θ_* as a random variable θ and computes $\mathbb{E}(\theta|Y_0 = y_0)$ as an estimator of θ_*. The key idea of the proposed method is to express the conditional expectation as a weighted sum of reverse conditional probabilities using Bayes' rule and to compute the reverse conditional probabilities using simulation. Our numerical results reveal that the proposed estimator is computed within a few seconds and successfully converges to the true value of $\mathbb{E}(\theta|Y_0 = y_0)$ as the computer time allocated to the simulation increases.

Future research topics include (1) extending the proposed method to the case where θ_* is multidimensional and Θ is continuous, (2) extending the proposed method to the case where we wish to estimate $\mathbb{E}(\theta|Y_0 = y_0, Y_{-1} = y_{-1}, \cdots, Y_{-m} = y_{-m})$ for some $m > 1$, and (3) justifying the proposed method rigorously by establishing the consistency of the proposed method.

References

Basawa, I. V. and Prakasa Rao, B. L. S. (1980). *Statistical Inference for Stochastic Processes.* Academic Press, New York.

Beadle, E. R. and Djurić, P. M. (1997). A fast weighted Bayesian bootstrap filter for nonlinear model state esstimation. *IEEE Trans. Aerosp. Electron. Syst.*, 33(1):338–343.

Bhada, A. and Ionides, E. L. (2014). Adaptive particle allocation in iterated sequential Monte Carlo via approximating meta-models. *Stat. Comput.*, pages 1–15.

Carpenter, J., Clifford, P., and Fearnhead, P. (1999). An improved particle filter for non-linear problems. *IEE Proceedings Radar, Sonar and Navigation*, 146(1):2–7.

Casella, G. and Robert, C. P. (1996). Rao-Blackwellisation of sampling schemes. *Biometrika*, 83(1):81–94.

Chen, H. and Yao, D. D. (2001). *Fundamentals of Queueing Networks: Performance, Asymptotics, and Optimization*. Springer, New York.

Doucet, A., Briers, M., and Senecal, S. (2006). Efficient block sampling strategies for sequential Monte Carlo. *J. Comput. Graph. Statist.*, 15(3):693–711.

Doucet, A. and Johansen, A. M. (2011). A tutorial on particle filtering and smoothing: Fifteen years later. In *Oxford Handbook of Nonlinear Filtering*. Oxford University Press.

Durrett, R. (2011). *Essentials of Stochastic Processes*. Springer, New York, NY.

Fox, D. (2003). Adapting the sample size in particle filters through KLD-sampling. *International Journal of Robotics Research*, 22(12):985–1004.

Gilks, W. R. and Berzuini, C. (2001). Following a moving target– Monte Carlo inference for dynamic Bayesian models. *J. R. Statist. Soc. B*, pages 127–146.

Jacob, P. E., Lindsten, F., and Schön, T. B. (2019). Smoothing with couplings of conditional particle filters. *J. Amer. Statist. Assoc.*, pages 1–20.

Kantas, N., Doucet, A., Singh, S. S., and Maciejowski, J. M. (2009). An overview of sequential Monte Carlo methods for parameter estimation in general state–space models. In *IFAC Proceedings*, pages 774–785.

Kitagawa, G. (1996). Montal Carlo filter and smoother for non-Gaussian nonlinear state space models. *J. Comput. Graph. Stat.*, 5(1):1–25.

Künsch, H. R. (2005). Recursive Monte Carlo filters: Algorithms and theoretical analysis. *Ann. Statist.*, 33(5):1983–2021.

Leippold, M. and Yang, H. (2019). Particle filtering, learning, and smoothing for mixed-frequency state-space models. *Econom. Stat.*, 12:25–41.

Liu, J. and Chen, R. (1998). Sequential Monte-Carlo methods for dynamic systems. *J. Amer. Statist, Assoc.*, 93(443):1032–1044.

Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag, New York.

Liu, J. S. and Chen, R. (1995). Blind deconvolution via sequential imputation. *J. Amer. Statist. Assoc.*, pages 567–576.

Nemeth, C., Fearnhead, P., and Mihaylova, L. (2014). Sequential Monte Carlo methods for state and parameter estimation in abruptly changing environments. *IEEE Trans. Signal Process.*, 62:1245–1255.

Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: Auxiliary particle filter. *J. Amer. Statist. Assoc.*, pages 590–599.

Yang, B., Stroud, J. R., and Huerta, G. (2018). Sequential Monte Carlo smoothing with parameter estimation. *Bayesian Anal.*, 13:1137–1161.