Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Simulation-Based Prediction

Eunji Lim, Peter W. Glynn

To cite this article:
Eunji Lim, Peter W. Glynn (2022) Simulation-Based Prediction. Operations Research

Published online in Articles in Advance 01 Feb 2022

. https://doi.org/10.1287/opre.2021.2229

http://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Articles in Advance, pp. 1–14
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

Simulation-Based Prediction

Eunji Lim,[a] **Peter W. Glynn**[b]

[a] Decision Sciences and Marketing, Adelphi University, Garden City, New York 11530; [b] Management Science and Engineering, Stanford University, Stanford, California 94305

Contact: elim@adelphi.edu, https://orcid.org/0000-0003-1008-7050 (EL); glynn@stanford.edu, https://orcid.org/0000-0003-1370-6638 (PWG)

Received: September 21, 2019
Revised: November 8, 2020; June 15, 2021; August 28, 2021
Accepted: October 13, 2021
Published Online in Articles in Advance: February 1, 2022

Area of Review: Simulation

https://doi.org/10.1287/opre.2021.2229

Abstract. This paper is concerned with the use of simulation in computing predictors in settings in which real-world observations are collected. A major challenge is that the state description underlying the simulation will typically include information that is not observed in the real system. This makes it challenging to initialize simulations that are aligned with the most recent observation collected in the real-world system, especially when the simulation does not visit the most recently observed value frequently. Our estimation methodology involves the use of "splitting," so that multiple simulations are launched from states that are closely aligned with the most recently collected real-world observation. We provide estimators both in the setting that the observed real-world values are discrete and are continuous, with kernel smoothing methods being systematically exploited in the continuous setting.

Supplemental Material: The e-companion is available at https://doi.org/10.1287/opre.2021.2229.

Keywords: prediction • initialization • splitting • Monte Carlo

1. Introduction

In many application settings, one has an interest in computing a prediction of future performance, based on the currently observed state of the system. Such problems arise naturally in various economic and operational settings, including supply chains, inventory systems, call centers, and manufacturing facilities. In fact, our interest in this problem was motivated within the setting of a container management system run by a major automobile manufacturer. The containers play a critical role in shipping parts from one manufacturing plant to another, and a shortage of containers at a plant can lead to a plant shutdown. In that setting, there is an interest in predicting future container shortages, based on information about the current locations of containers. A key complication is that the exact current location of each container is not observed, since individual containers often are tracked only when they move through certain transshipment points. So, the observable data, in this setting, are not aligned with the full "state" necessary in order to initialize a dynamic simulation of the container system. We note that such container shortages would be of interest days or weeks into the future, so that significant computing time would be available to run the simulations associated with our contribution.

In this paper, we develop a framework and algorithms for solving prediction problems of this type. In particular, we consider the problem of predicting the real-valued random variable (rv) $\widetilde{Z}(s+t)$ for $t \geq 0$, based on having observed the S'-valued process $(\widetilde{Y}(u) : 0 \leq u \leq s)$, where $S' \subset \mathbb{R}^d$. In the container management system that motivated this paper, $\widetilde{Y}(u)$ is the numbers of containers at time u in the transshipment points where containers are tracked, s is the current time, $s + t$ is the future point in time at which we are interested in making a prediction, and $\widetilde{Z}(s+t)$ is what we wish to predict, such as the shortage in the number of containers at one of the manufacturing plants at time $s + t$. We simplify the problem by basing our prediction on $\widetilde{Y}(s)$ only, leaving the question of utilizing the full path history $(\widetilde{Y}(u) : 0 \leq u \leq s)$ to future work. Such a "full history" calculation would require solving the filtering problem for general discrete-event simulations, a problem beyond the scope of this paper. For our current purposes, our goal is to minimize the mean square prediction error between $\widetilde{Z}(s+t)$ and a predictor of $\widetilde{Z}(s+t)$ given $\widetilde{Y}(s)$, so our desired predictor is the conditional expectation of $\widetilde{Z}(s+t)$ given $\widetilde{Y}(s)$, that is, $\mathbb{E}[\widetilde{Z}(s+t) \mid \widetilde{Y}(s)]$ (see, e.g., section 9.4 on p. 85 of Williams 1991 or p. 468 of Karlin and Taylor 1975). Of course, we have in mind settings where this conditional expectation is not available in closed form and must be computed numerically (e.g., by simulation).

Recall that the observable $\widetilde{Y}(s)$ is typically not the state variable necessary in order to initialize a simulation of the system over $[s, s+t]$. We presume that there exists an underlying S-valued Markov process $\widetilde{X} = (\widetilde{X}(u) : u \geq 0)$ for which $\widetilde{Z}(u) = f(\widetilde{X}(u))$ and $\widetilde{Y}(u) = \tau(\widetilde{X}(u))$, for suitably chosen (known) measurable

functions $f : S \to \mathbb{R}$ and $\tau : S \to \mathbb{R}^d$. In the container management system where we wish to predict the number of containers in shortage, $\widetilde{X}(u)$ includes the numbers of containers at each of the manufacturing plants at time u, along with other state variables, such as the elapsed interarrival times or elapsed service times. For $x \in S$, $f(x)$ is the number of orders waiting for available containers (assuming that each order requires exactly one container for its delivery), and $\tau(x)$ consists of all observable state variables, such as the number of containers at the observed transshipment points. See Section 6.2.2 for a more detailed description of the state variables of X, f, and τ in the setting of an example.

We can then represent the conditional expectation of interest as $k(\mathbf{y})$ when $\widetilde{Y}(s) = \mathbf{y}$ is the observed value at time s, where

$$k(\mathbf{y}) \triangleq \mathbb{E}[\widetilde{Z}(s+t) \mid \widetilde{Y}(s) = \mathbf{y}] = \int_S \mathbb{E}[f(\widetilde{X}(s+t))|\widetilde{X}(s) = \mathbf{x}] \times \mathbb{P}(\widetilde{X}(s) \in d\mathbf{x}|\widetilde{Y}(s) = \mathbf{y}).$$

For each $\mathbf{x} \in S$, the inner expectation can be estimated via standard simulation methods as a finite-horizon (terminating) simulation initialized from $\mathbf{x}$. The challenge, of course, is performing the outer integration over S and implementing the entire calculation in an integrated and efficient manner. In particular, the outer integration involves sampling $\widetilde{X}(s)$, conditional on $\widetilde{Y}(s) = \mathbf{y}$. Given that the observed system has likely been in operation for a long period of time, the process $\widetilde{Y}$ can be assumed to be in a steady state. It follows that generating $\widetilde{X}(s)$ given $\widetilde{Y}(s) = \mathbf{y}$ involves all the complications that typically arise in the context of steady-state simulation, including that of assessing the variance in the context of an autocorrelated process. Another issue is that $\mathbf{y}$ may be infrequently visited by $\widetilde{Y}$, and this creates an algorithmic challenge.

1.1. Proposed Methodology

To develop an efficient way to sample $\widetilde{X}(s)$ conditional on $\widetilde{Y}(s) = \mathbf{y}$, we recognize that if $\widetilde{X}$ has been simulated for a long period of time, then $\widetilde{X}$ can be viewed as a strictly stationary Markov process. Under this assumption, one obvious way to sample $\widetilde{X}(s)$ conditional on $\widetilde{Y}(s) = \mathbf{y}$ is by simulating $(\widetilde{X}(u) : u \geq 0)$ and sampling $\widetilde{X}(u)$ only when $\widetilde{Y}(u) = \mathbf{y}$. Two main challenges become apparent from a computational point of view.

1. It may require a long period of time until $(\widetilde{X}(u) : u \geq 0)$ gets close to its stationary distribution. This arises because the simulation run is typically initialized using a nonequilibrium distribution.

2. Even if the simulation gets close to its stationary distribution, $\widetilde{Y}(u)$ may hit $\mathbf{y}$ infrequently, especially when S' has a large number of states or is a continuous state space.

To tackle the first challenge, we generate a single long simulation run of $(\widetilde{X}(u) : u \geq 0)$ rather than multiple independent replications, and we use the simulated values of $\widetilde{X}(u)$ for all the u values where $\widetilde{Y}(u) = \mathbf{y}$. Typically, a single run has computational advantages over multiple independent runs when one wishes to generate samples from the stationary distribution because each independent run has an initial transient phase prior to a steady state (see, e.g., p. 3 of Nelson 2016). It should be noted that, in this single run, we collect $\widetilde{X}(u)$ and $\widetilde{Z}(u+t)$ whenever $\widetilde{Y}(u) = \mathbf{y}$ for $u \geq 0$, so our sampled data are dependent.

To tackle the second challenge, we use a couple of strategies. First, we use "splitting," in which whenever $\widetilde{Y}(u)$ hits $\mathbf{y}$, we conduct multiple independent simulation runs initialized from $\widetilde{X}(u)$, thereby generating multiple samples of $\widetilde{Z}(u+t)$ that are independent conditional on $\widetilde{X}(u)$ (see Figure 1 for a graphical representation of this idea). In the second strategy, we use $\widetilde{Z}(u+t)$, even if

Figure 1. Representation of the Splitting Strategy

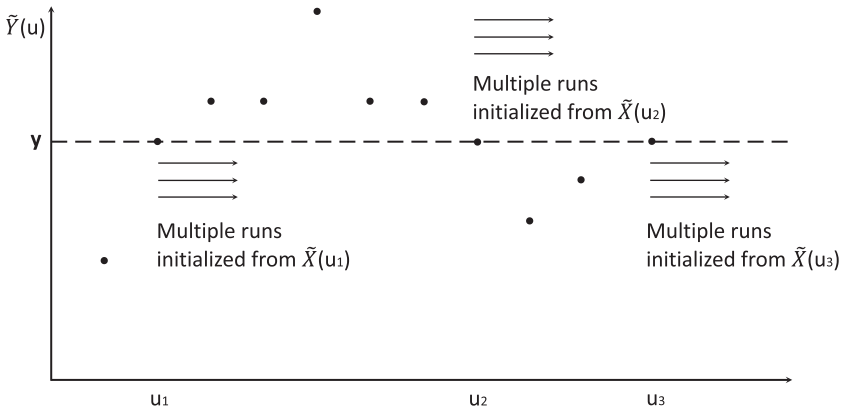

Notes. The x-axis is $u \geq 0$. The dots are $\widetilde{Y}(u)$. Whenever $\widetilde{Y}(u)$ hits $\mathbf{y}$, we generate multiple runs of $\widetilde{X}$ initialized from the corresponding value of $\widetilde{X}(u)$.

$\widetilde{Y}(u) \neq \mathbf{y}$. Specifically, we assign a weight to $\widetilde{Z}(u+t)$ so that more weight is applied when $\widetilde{Y}(u)$ is closer to $\mathbf{y}$. We then compute the weighted average of $\widetilde{Z}(u+t)$ for $u \geq 0$ as an estimator of $k(\mathbf{y})$. This strategy is closely related to the ideas arising in kernel regression estimation.

Our proposed estimator of $k(\mathbf{y})$ in the discrete setting, $k_n(\mathbf{y})$, uses the splitting idea, while an alternative estimator, $k_n^*(\mathbf{y})$, uses both the splitting idea and the kernel regression idea. Theorem 3 in Section 4 and Lemma 6 in the e-companion to this paper state that when the discrete-time version of $\widetilde{X}$ is a Harris ergodic Markov chain satisfying certain moment and mixing conditions, regardless of how the simulation is initialized, both $\sqrt{n}(k_n(\mathbf{y}) - k(\mathbf{y}))$ and $\sqrt{n}(k_n^*(\mathbf{y}) - k(\mathbf{y}))$ converge weakly to a normal distribution as $n \to \infty$, and their asymptotic variance is a decreasing function of the split factor, which is the number of multiple samples of $\widetilde{Z}(u+t)$.

1.2. Literature Review

According to whether $t > 0$, $t = 0$, or $t < 0$, the problem of estimating $\mathbb{E}[\widetilde{Z}(s+t)|\widetilde{Y}(s) = \mathbf{y}]$ has been referred to as *prediction*, *filtering*, and *smoothing*, respectively, in the statistics literature. When $(\widetilde{Z}(u) : u \geq 0)$ is observable, the problem of predicting $\mathbb{E}[\widetilde{Z}(s+t)|\widetilde{Y}(s) = \mathbf{y}]$ is closely related to the regression estimation problem, since $k(\mathbf{y}) = \mathbb{E}[\widetilde{Z}(s+t)|\widetilde{Y}(s) = \mathbf{y}]$ can be treated as a regression function in $\mathbf{y}$, and $k(\mathbf{y})$ can be estimated using the observational data $((\widetilde{Y}(u), \widetilde{Z}(u)) : 0 \leq u \leq s)$. Numerous parametric and nonparametric regression approaches have been proposed and studied extensively in the literature (see, e.g., Györfi et al. 2002, for a comprehensive survey). On the other hand, when $(\widetilde{Z}(u) : u \geq 0)$ is not observable, simulation must be conducted to obtain simulated values of $(\widetilde{Z}(u) : u \geq 0)$ from simulated paths of the underlying Markov process $\widetilde{X}$. In order to conduct the simulation, the need for simulating from $\mathbb{P}(\widetilde{X}(s) \in \cdot\,|\widetilde{Y}(s) = \mathbf{y})$ arises.

One closely related question is the problem of estimating $\mathbb{P}(\widetilde{X}(s) \in \cdot\,|\widetilde{Y}(s) = \mathbf{y}_s, \widetilde{Y}(s-1) = \mathbf{y}_{s-1}, \ldots, \widetilde{Y}(s-r) = \mathbf{y}_{s-r})$ for some $r > 0$, which is answered by numerous sequential Monte Carlo methods. However, the sequential Monte Carlo methods assume that one is able to draw samples from $\mathbb{P}(\widetilde{X}(s-r) \in \cdot\,|\widetilde{Y}(s-r) = \mathbf{y}_{s-r})$ and update the samples iteratively to get an approximation to $\mathbb{P}(\widetilde{X}(s) \in \cdot\,|\widetilde{Y}(s) = \mathbf{y}_s, \widetilde{Y}(s-1) = \mathbf{y}_{s-1}, \ldots, \widetilde{Y}(s-r) = \mathbf{y}_{s-r})$, so the question of how to sample from $\mathbb{P}(\widetilde{X}(s-r) \in \cdot\,|\widetilde{Y}(s-r) = \mathbf{y}_{s-r})$ remains unanswered (see, e.g., Liu and West 2001). Therefore, this paper is the first to address the need for simulating from $\mathbb{P}(\widetilde{X}(s) \in \cdot\,|\widetilde{Y}(s) = \mathbf{y})$ when estimating $k(\mathbf{y})$. We view the introduction of this class of problems to the research community as a major contribution of this paper.

Another issue with the sequential Monte Carlo algorithms is that their typical computational burden is prohibitively heavy (see, e.g., Pitt and Shephard 1999 and Doucet and Johansen 2009). This paper solves the issue of the computational burden by utilizing the idea of splitting in conjunction with kernel estimation. The idea of splitting has received a great deal of interest in rare event simulation (see p. 127 of Asmussen and Glynn 2007 and the references therein). However, most splitting-based methods terminate the simulation whenever $\widetilde{Y}(u) = \mathbf{y}$ occurs. Our approach applies the splitting idea every time $\widetilde{Y}(u)$ hits $\mathbf{y}$ in a long single run of $\widetilde{X}$, so the successive events $\widetilde{Y}(u) = \mathbf{y}$ are dependent on one another. Thus, this paper is the first to apply the splitting idea to successive events $\widetilde{Y}(s) = \mathbf{y}$ that are dependent. As our theoretical and empirical studies in Sections 4 and 6 suggest, the splitting idea plays an important role in reducing the asymptotic variance of the proposed estimators.

The focus of this paper is placed on the case where $\mathbf{y}$ lives in a large discrete space (such as the case of the container management system), so one naturally prefers a method that does not assume any functional form on $k(\cdot)$. If, however, one can make some assumptions on $k(\cdot)$, such as the continuity of $k(\cdot)$ in the continuous variable $\mathbf{y}$, then more sophisticated methods can be proposed. For example, Hong and Jiang (2019) introduced a new framework, called *offline-simulation online-application*, in which one generates simulated data for $k(\mathbf{z})$ offline for various values of $\mathbf{z}$, builds a metamodel for $k(\cdot)$ with the simulated data, and uses the metamodel to estimate $k(\mathbf{y})$ or to solve real-time decision problems. This approach is particularly suitable when one can assume that $\mathbf{y}$ resides in a continuous space and $k(\cdot)$ is continuous over its domain, so various surface fitting methods can be used to build such a metamodel. Related numerical results can be found in Section 6.1.4.

In Section EC.5 of the e-companion to this paper, we discuss the problem of estimating unknown parameters of a stochastic model based on observed, real-world data, which has gained increasing interest in the simulation community (see Peng et al. 2020 for a recent development).

1.3. Organization of This Paper

The rest of this paper is organized as follows. Section 2 introduces some notation and definitions. In Section 3, we briefly review the development of predictors that do not rely on simulation and that use only the observed real-world data that are available. Section 4 contains our main results and describes a set of simulation-based predictors. All of our simulation-based predictors take advantage of splitting, so that multiple conditionally independent replications of $\widetilde{Z}(u+t)$ are generated from a single $\widetilde{X}(u)$ (whenever the corresponding $\widetilde{Y}(u)$ is sufficiently close to $\mathbf{y}$). Section 5 provides a brief discussion of a couple of extensions of the methodology presented in

Section 4. These include simulation-based prediction when $\widetilde{Y}(s)$ is discrete but takes on a very large number of possible values, and the construction of confidence intervals for our simulation-based predictors. In Section 6, we present the numerical performance of our proposed estimators in the container management example. Concluding remarks are included in Section 7.

2. Notation and Definitions

We let $\Rightarrow$ and $\xrightarrow{P}$ denote weak convergence and convergence in probability, respectively. By $N(0, 1)$, we denote the normal random variable with a mean of 0 and a variance of 1.

Suppose that $V = (V_i : i = 0, 1, \cdots)$ is a sequence of random vectors, each of which takes values in $S \subset \mathbb{R}^p$. For $0 \le i \le j \le \infty$, let $\mathcal{F}_i^j$ denote the σ-field generated by $V_i, \ldots, V_j$. For each $n \ge 1$, we define the following strong mixing coefficients:

$$\alpha_n = \sup_{k=0,1,\cdots} \sup\{|\mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B)| : A \in \mathcal{F}_0^k, B \in \mathcal{F}_{k+n}^\infty\}.$$

The sequence V is said to be *strongly mixing* (or *α-mixing*) if $\alpha_n \to 0$ as $n \to \infty$.

For $\mathbf{x} = (x_1, \ldots, x_d) \in \mathbb{R}^d$, we let $\|\mathbf{x}\| = (x_1^2 + \cdots + x_d^2)^{1/2}$. For any sequences $(f_n : n = 1, 2, \cdots)$ and $(g_n : n = 1, 2, \cdots)$ of real numbers, we write $f_n \sim g_n$ if $f_n/g_n \to 1$ as $n \to \infty$, $f_n = o(g_n)$ if $f_n/g_n \to 0$ as $n \to \infty$, and $f_n = O(g_n)$ if there exist positive real number M and a positive integer N such that $|f_n| \le M g_n$ for all $n \ge N$. For $x \in \mathbb{R}$, $\lfloor x \rfloor$ denotes the largest integer that is less than or equal to x.

For a sequence of random variables $(Z_n : n \ge 1)$ and a sequence of positive real numbers $(a_n : n \ge 1)$, we say that $Z_n = O_p(a_n)$ as $n \to \infty$ if, for any $\epsilon > 0$, there exist positive constants c_0 and n_0 such that $\mathbb{P}(|Z_n/a_n| > c_0) < \epsilon$ for all $n \ge n_0$. We also say that $Z_n = o_p(a_n)$ as $n \to \infty$ if Z_n/a_n converges in probability to 0.

A discrete-time Markov chain $V = (V_i : i = 0, 1, \ldots)$ is called *ergodic* if it is positive Harris recurrent and aperiodic.

3. Prediction Based on Observed Statistical Data Only

To help understand the benefit of using simulation for such prediction problems, we first consider the case in which no simulation is used, so that only observed statistical data are utilized in estimating the quantity $k(\mathbf{y})$. In particular, we assume here that $((\widetilde{Z}(u), \widetilde{Y}(u)) : 0 \le u \le s)$ is (statistical) data observed from the real-world setting of interest.

Let $w = s - \lfloor s \rfloor$. For $i = 0, 1, \cdots$, put $\widetilde{X}_i = \widetilde{X}(i + w)$, $\widetilde{Y}_i = \widetilde{Y}(i + w)$, and $\widetilde{Z}_i = \widetilde{Z}(i + w + t)$. Throughout this paper, we will use the following assumption.

A0. *We have that $(\widetilde{X}_i : i = 0, 1, \cdots)$ is strictly stationary.*

This assumption is less restrictive than assuming that $\widetilde{X}$ is strictly stationary. For example, A0 covers the case in which $\widetilde{X}$ is a Markov process with periodic nonhomogeneous transition probabilities having period 1. Such periodic nonhomogeneity is needed to cover the many operations management settings in which time-of-day or day-of-week effects have a significant impact on the system dynamics.

For the analysis in this section, set $r - 1 = \lfloor s - w - t \rfloor$. We have set $r - 1$ this way so that when $i = r - 1$, $i + w + t$ does not exceed s.

3.1. Discrete Case

Suppose that S' is either finite or countably infinite and that $\widetilde{Y}_0$ is a discrete S'-valued rv. Assume that $p(\mathbf{y}) \triangleq \mathbb{P}(\widetilde{Y}_0 = \mathbf{y}) > 0$. In that case, we may consider the estimator

$$\widetilde{k}_r(\mathbf{y}) = \frac{\sum_{i=0}^{r-1} \widetilde{Z}_i I(\widetilde{Y}_i = \mathbf{y})}{\sum_{i=0}^{r-1} I(\widetilde{Y}_i = \mathbf{y})}.$$

Below are the key assumptions that we will use when analyzing the asymptotic behavior of $\widetilde{k}_r(\mathbf{y})$.

A1.

(a) *There exists a positive real number σ_1 such that*

$$r^{-1/2} \sum_{i=0}^{r-1} (\widetilde{Z}_i - k(\mathbf{y})) I(\widetilde{Y}_i = \mathbf{y}) \Rightarrow \sigma_1 N(0, 1) \tag{3.1}$$

as $r \to \infty$.

(b) *We have that*

$$\frac{1}{r} \sum_{i=0}^{r-1} I(\widetilde{Y}_i = \mathbf{y}) \to p(\mathbf{y}) \tag{3.2}$$

almost surely (a.s.) as $r \to \infty$.

There exist a number of sets of conditions under which A1 holds, so we present one of them here.

B1. *We have that $(\widetilde{X}_i : i = 0, 1, \ldots)$ is strongly mixing with the strong mixing coefficients $(\alpha_n : n = 1, 2, \ldots)$. There exists a positive real number δ such that*

(a) $\mathbb{E}[|\widetilde{Z}_0|^{2+\delta}] < \infty$, *and*

(b) $\sum_{n=1}^\infty \alpha_n^{\delta/(2+\delta)} < \infty$.

Theorem 1 identifies the limiting behavior of $\widetilde{k}_r(\mathbf{y})$ under A0 and B1. The proof of Theorem 1 is provided in Section EC.1 of the e-companion to this paper.

Theorem 1. *Under A0 and B1,*

$$\sigma^2 \triangleq \mathrm{var}[(\widetilde{Z}_0 - k(\mathbf{y})) I(\widetilde{Y}_0 = \mathbf{y})] + 2 \sum_{i=1}^\infty \mathrm{cov}[(\widetilde{Z}_0 - k(\mathbf{y})) I(\widetilde{Y}_0 = \mathbf{y}), (\widetilde{Z}_i - k(\mathbf{y})) I(\widetilde{Y}_i = \mathbf{y})]$$

is convergent and A1(b) holds. Furthermore, if $\sigma^2 > 0$, then A1(a) holds with $\sigma_1 = \sigma$, and

$$r^{1/2}\left(\widetilde{k}_r(\mathbf{y}) - k(\mathbf{y})\right) \Rightarrow \frac{\sigma_1}{p(\mathbf{y})} N(0, 1) \tag{3.3}$$

as $r \to \infty$.

One way to establish B1(b) is by requiring $(\widetilde{X}_i : i = 0,1,\cdots)$ to be geometrically ergodic, since any aperiodic, Harris recurrent, and geometrically ergodic Markov chain has exponentially decaying strong mixing coefficients (see, e.g., Nummelin and Tuominen 1982 and p. 199 of Rosenblatt 1971). Hence, we can establish the following corollary.

Corollary 1. *Suppose that $(\widetilde{X}_i : i = 0,1,\cdots)$ is an aperiodic, Harris recurrent, and geometrically ergodic Markov chain. Assume that $\widetilde{X}_0$ is initialized with its stationary distribution. Assume also that $\mathbb{E}\,|\widetilde{Z}_0|^{2+\delta} < \infty$ for some positive real number δ. Then,*

$$\sigma_1^2 \triangleq \mathrm{var}[(\widetilde{Z}_0 - k(\mathbf{y}))I(\widetilde{Y}_0 = \mathbf{y})] + 2\sum_{i=1}^{\infty} \mathrm{cov}[(\widetilde{Z}_0 - k(\mathbf{y}))I(\widetilde{Y}_0 = \mathbf{y}), (\widetilde{Z}_i - k(\mathbf{y}))I(\widetilde{Y}_i = \mathbf{y})]$$

is convergent, and if $\sigma_1^2 > 0$,

$$r^{1/2}\left(\widetilde{k}_r(\mathbf{y}) - k(\mathbf{y})\right) \Rightarrow \frac{\sigma_1}{p(\mathbf{y})} N(0,1) \tag{3.4}$$

as $r \to \infty$.

Remark 1. Many Markov chains are geometrically ergodic. For example, if a positive Harris recurrent and aperiodic Markov chain satisfies the Doeblin condition, then it is geometrically ergodic (see, e.g., Doob 1953). Also, any irreducible aperiodic finite-state Markov chain is geometrically ergodic (see p. 18 of Asmussen 2003).

3.2. Continuous Case

We next turn to the setting in which $\mathbb{P}(\widetilde{Y}_0 = \mathbf{y}') = 0$ for $\mathbf{y}' \in S'$, as occurs when $\widetilde{Y}(s)$ is a continuous S'-valued rv. In particular, suppose that $\widetilde{Y}_0$ is an $\mathbb{R}^d$-valued rv with density $g(\cdot)$ for which $g(\mathbf{y}) > 0$. Let $(a_n : n = 1,2,\cdots)$ be a "bandwidth" sequence decreasing to 0, and put $\lambda_{ir}(\mathbf{y}) = \lambda((\mathbf{y} - \widetilde{Y}_i)/a_r)$ for some "kernel" function $\lambda : \mathbb{R}^d \to \mathbb{R}$. In order that the number of samples in a neighborhood of $\mathbf{y}$ having radius a_n tends to infinity, we assume that $na_n^d \to \infty$ as $n \to \infty$. The estimator

$$\widetilde{k}_r'(\mathbf{y}) = \frac{\sum_{i=0}^{r-1} \widetilde{Z}_i \lambda_{ir}(\mathbf{y})}{\sum_{i=0}^{r-1} \lambda_{ir}(\mathbf{y})}$$

is then a natural choice of "kernel-based" estimator for $k(\mathbf{y})$. This is known in the literature as the *Nadaraya-Watson estimator* for the "regression function" $k(\cdot)$ (see, e.g., Györfi et al. 2002 and Abdellah et al. 2018).

There exist numerous results establishing central limit theorems (CLTs) for $\widetilde{k}_r'(\mathbf{y})$. We introduce one of them here under the following set of assumptions.

B2.

(a) *We have that $a_n = O(n^{-\gamma})$ as $n \to \infty$ for some positive constant γ, and $na_n^d \to \infty$ as $n \to \infty$.*

(b) *We have that $\lambda : \mathbb{R}^d \to \mathbb{R}$ is symmetrical with respect to the origin, that is, $\lambda(\mathbf{z}) = \lambda(\mathbf{w})$ if $\|\mathbf{z}\| = \|\mathbf{w}\|$ for $\mathbf{z}, \mathbf{w} \in \mathbb{R}^d$. Define $f_\lambda : [0,\infty) \to \mathbb{R}$ by $f_\lambda(\|\mathbf{z}\|) = \lambda(\mathbf{z})$ for $\mathbf{z} \in \mathbb{R}^d$; λ is a bounded nonnegative function with support $\{\mathbf{z} \in \mathbb{R}^d : \|\mathbf{z}\| \le 1\}$.*

(c) *There exist positive real numbers β and c_1 such that*

$$|k(\mathbf{z}_1) - k(\mathbf{z}_2)| \le c_1 \|\mathbf{z}_1 - \mathbf{z}_2\|^\beta$$

for all $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^d$ and $1 - 2\gamma\beta - \gamma d < 0$.

(d) *We have that $g_2 : \mathbb{R}^d \to \mathbb{R}$ is continuous in some neighborhood of $\mathbf{y}$, where $g_2(\mathbf{z}) = \mathrm{var}[\widetilde{Z}_0 | \widetilde{Y}_0 = \mathbf{z}]$ for $\mathbf{z} \in \mathbb{R}^d$.*

(e) *Assume that $\mathbb{E}\,|\widetilde{Z}_0|^\nu < \infty$ for some $\nu > 2$. Let $g_\nu : \mathbb{R}^d \to \mathbb{R}$ be defined by*

$$g_\nu(\mathbf{z}) = \mathbb{E}[|\widetilde{Z}_0 - k(\mathbf{y})|^\nu \mid \widetilde{Y}_0 = \mathbf{z}]$$

for $\mathbf{z} \in \mathbb{R}^d$; g_ν is continuous in some neighborhood of $\mathbf{y}$.

(f) *Define $g_{ij}(\mathbf{z}_1, \mathbf{z}_2) = \mathbb{E}[(\widetilde{Z}_i - k(\mathbf{y}))(\widetilde{Z}_j - k(\mathbf{y})) \mid \widetilde{Y}_i = \mathbf{z}_1, \widetilde{Y}_j = \mathbf{z}_2]$ for $i \ne j$, $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^d$. Assume that g_{ij} is bounded in some neighborhood of $(\mathbf{y}, \mathbf{y})$ uniformly in i and j; that is, there exist positive constants h and c_2 such that*

$$\sup_{i \ne j} \sup_{\|\mathbf{z}_1 - \mathbf{y}\| \le h, \|\mathbf{z}_2 - \mathbf{y}\| \le h} |g_{ij}(\mathbf{z}_1, \mathbf{z}_2)| \le c_2.$$

(g) *There exist functions $\psi : \mathbb{R} \to \mathbb{R}$, satisfying $\lim_{h \to 0} \psi(h) = 0$, and $f_1 : \mathbb{R}^d \to \mathbb{R}$ such that, $\sup_{i \ne j} \mathbb{P}(\|\widetilde{Y}_i - \mathbf{y}\| \le h, \|\widetilde{Y}_j - \mathbf{y}\| \le h) \le \psi(h) f_1(\mathbf{y})$ as $h \to 0$. We also assume that $\psi(h)/h^{2d}$ is bounded.*

(h) *We have that $(\widetilde{X}_i : i = 0,1,\cdots)$ is strongly mixing with the strong mixing coefficients $(\alpha_n : n = 1,2,\cdots)$ satisfying*

$$\sum_{n=1}^{\infty} n^\delta \alpha_n^{1-2/\kappa} < \infty$$

for some $\kappa > 2$ and $\delta > 1 - 2/\kappa$.

The following result is due to equation (2.11) and corollary 2 of Masry (2005).

Theorem 2. *Under A0 and B2,*

$$(ra_r^d)^{1/2}\left(\widetilde{k}_r'(\mathbf{y}) - k(\mathbf{y})\right) \Rightarrow \sigma_2 N(0,1)$$

as $r \to \infty$, where

$$\sigma_2^2 = \mathrm{var}[\widetilde{Z}_0 \mid \widetilde{Y}_0 = \mathbf{y}] \frac{\int_0^1 \{f_\lambda(x)\}^2 dx}{g(\mathbf{y})\left(\int_0^1 f_\lambda(x) dx\right)^2}.$$

Remark 2. One way to establish B2(h) is by requiring one of the following conditions:

1. α_n decays exponentially fast; that is, $\alpha_n = O(e^{-\alpha n})$ as $n \to \infty$ for some $\alpha > 0$.

2. $\alpha_n = O(n^{-\alpha})$ as $n \to \infty$ for some positive constant α satisfying

$$\alpha > \max\left\{\frac{2}{\gamma d} - 1, (2 - 2/\nu)/(1 - 2/\nu)\right\}.$$

Thus, as in Corollary 1, B2(h) is satisfied if $(\widetilde{X}_i : i = 0,1,\cdots)$ is an aperiodic, Harris recurrent, and

geometrically ergodic Markov chain. (For more detailed explanations on B2(h), see remark 3 on p. 163 of Masry 2005.)

Remark 3. For a discussion on how to choose λ, see section 5 of Bierens (1987).

4. Simulation-Based Predictors

We now turn to the development of simulation-based predictors for $k(\mathbf{y})$. As we will see, use of simulation can significantly improve the quality of our estimator for $k(\mathbf{y})$, relative to those introduced in Section 3. In particular, the sample sizes available in the simulation setting have the potential to be vastly larger than those based on observed statistical data, where it is likely that the parameter s will often be small or of only moderate size. Furthermore, it may be that the quantity $\widetilde{Z}(s+t)$ to be predicted is not observed in the real-world system, so that the estimators of Section 3 are not applicable. This would arise, for example, in the container setting when one wishes to predict container shortages in a part of the system in which limited monitoring is available.

In the simulation setting, we simulate the processes $\widetilde{X}, \widetilde{Y}$, and $\widetilde{Z}$. To differentiate our simulation-based data from the real-world observed statistical data, we denote the simulation counterparts of $\widetilde{X}, \widetilde{Y}$, and $\widetilde{Z}$ by X, Y, and Z, respectively. The process X will share the same transition probabilities as $\widetilde{X}$ but typically will be initialized differently. In particular, X will not exhibit the periodic stationarity of $\widetilde{X}$ because the simulation will generally be initialized from a nonequilibrium distribution. However, under suitable regularity conditions on X, the marginal probabilities agree asymptotically in the sense that $(X_i, Y_i, Z_i) \Rightarrow (\widetilde{X}_0, \widetilde{Y}_0, \widetilde{Z}_0)$ as $i \to \infty$, where $X_i = X(i+w)$, $Y_i = Y(i+w)$, and $Z_i = Z(i+w+t)$ for $i = 0, 1, \cdots$. (As in Section 3, we let $w = s - \lfloor s \rfloor$.) As noted in the introduction, in the simulation setting, we must typically simulate the process X in order to generate Y and Z. An important idea that we will exploit algorithmically is that we have the ability to "save" the state X_i and to potentially generate multiple forward (conditionally independent) simulations of Z to time $i + w + t$ using X_i as the "initial condition" for these simulations. This application of "splitting" (see p. 127 of Asmussen and Glynn 2007) is especially useful in this prediction context.

To be specific, for $i = 0, 1, \cdots$ and $j = 1, 2, \cdots, m$, let $Z_{ij} = f(X_{ij}(i+w+t))$, where $(X_{ij}(u) : i+w \le u \le i+w+t)$ is a forward simulation of $\widetilde{X}$ from time $i + w$ to time $i + w + t$, using X_i as the initial condition at time $i + w$. Given X_i, the multiple forward simulations $(X_{ij}(u) : i+w \le u \le i+w+t)$ for $j = 1, 2, \cdots, m$ should be generated independently of each other, and of Z_i. We can express the independence mathematically as follows:

$$\mathbb{P}(Z_{ij} \in A_{ij}, i = 0, 1, \cdots, n, j = 1, 2, \cdots, m \mid X_0, \cdots, X_n) = \prod_{i=0}^{n} \prod_{j=1}^{m} \mathbb{P}(Z_i \in A_{ij} \mid X_0, \cdots, X_n) \tag{4.1}$$

for any measurable subset $A_{ij} \subseteq \mathbb{R}$. When computing our proposed estimators in Sections 4.1 and 4.2, one needs to generate the Z_{ij}'s only when $Y_i = \mathbf{y}$ or $\lambda((Y_i - \mathbf{y})/a_n)$ is positive.

Throughout this section, we let $\mu(\mathbf{x}) = \mathbb{E}[Z_0 \mid X_0 = \mathbf{x}]$ for $\mathbf{x} \in S$.

4.1. Discrete Case: Proposed Estimator $k_n(\mathbf{y})$

Suppose that S' is either finite or countably infinite, so $\widetilde{Y}_0$ is an S'-valued discrete rv. Assume that $p(\mathbf{y}) = \mathbb{P}(\widetilde{Y}_0 = \mathbf{y}) > 0$. After generating a single run $((X_i, Y_i, Z_i) : i = 0, 1, \cdots)$ along with the Z_{ij}'s, we estimate $k(\mathbf{y})$ via

$$k_n(\mathbf{y}) = \frac{\sum_{i=0}^{n-1} \left(\frac{1}{m} Z_i + \frac{1}{m} \sum_{j=1}^{m-1} Z_{ij} \right) I(Y_i = \mathbf{y})}{\sum_{i=0}^{n-1} I(Y_i = \mathbf{y})}$$

for $n = 1, 2, \cdots$.

To analyze the estimator $k_n(\mathbf{y})$, we use the following assumptions.

A2.

(a) *There exists a positive real number σ_3 such that, as $n \to \infty$,*

$$n^{-1/2} \sum_{i=0}^{n-1} \left(\frac{1}{m} (Z_i - \mu(X_i)) + (\mu(X_i) - k(\mathbf{y})) \right) I(Y_i = \mathbf{y}) \Rightarrow \sigma_3 N(0, 1). \tag{4.2}$$

(b) *There exists a positive real number σ_3' such that, conditional on X, as $n \to \infty$,*

$$n^{-1/2} \sum_{i=0}^{n-1} \left(\frac{1}{m-1} \sum_{j=1}^{m-1} (Z_{ij} - \mu(X_i)) I(Y_i = \mathbf{y}) \right) \Rightarrow \frac{\sigma_3'}{\sqrt{m-1}} N(0, 1). \tag{4.3}$$

(c) *We have that*

$$\frac{1}{n} \sum_{i=0}^{n-1} I(Y_i = \mathbf{y}) \to p(\mathbf{y}) \tag{4.4}$$

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

We now introduce a set of conditions under which A2 holds.

B3. *We have that $(\widetilde{X}_i : i = 0, 1, \ldots)$ is a strongly mixing sequence with the strong mixing coefficients $(\alpha_n : n = 1, 2, \ldots)$. There exists a positive real number δ such that*

(a) $\mathbb{E}[|\widetilde{Z}_0|^{2+\delta}] < \infty$,
(b) $\sum_{n=1}^{\infty} \alpha_n^{\delta/(2+\delta)} < \infty$,
(c) $\sup_{x\in S} \mathbb{E}[|\widetilde{Z}_0|^{2+\delta} | \widetilde{X}_0 = \mathbf{x}] < \infty$, *and* $\mathbb{E}[|\mathrm{var}[\widetilde{Z}_0|\widetilde{X}_0]|^{2+\delta}] < \infty$.

With B3 in place, we have the following theorem describing the behavior of $k_n(\mathbf{y})$ when the simulation budget n goes to infinity. The proof of Theorem 3 is provided in Section EC.1 of the e-companion to this paper.

Theorem 3. *Suppose that* $(\widetilde{X}_i : i = 0, 1, \dots)$ *is a Harris ergodic (i.e., aperiodic and positive recurrent) Markov chain. Let* $(X_i : i = 0, 1, \dots)$ *be a Markov chain with the same transition kernel as that of* $(\widetilde{X}_i : i = 0, 1, \dots)$ *and an arbitrary initial distribution. Let* $(Z_{ij} : i = 0, 1, \dots, j = 1, 2, \dots)$ *be a family of real-valued rv's satisfying* (4.1). *Let* $\widetilde{A}_i \triangleq (\frac{1}{m}(\widetilde{Z}_i - \mu(\widetilde{X}_i)) + (\mu(\widetilde{X}_i) - k(\mathbf{y})))I(\widetilde{Y}_i = \mathbf{y})$ *for* $i = 0, 1, \dots$. *Under* A0 *and* B3,

$$\sigma^2(m) \triangleq \mathrm{var}[\widetilde{A}_0] + 2\sum_{i=1}^{\infty} \mathrm{cov}[\widetilde{A}_0, \widetilde{A}_i]$$

is convergent, and items (b) and (c) of A2 *hold with* $\sigma_3' = \eta(\mathbf{y})\sqrt{p(\mathbf{y})}$, *where* $\eta(\mathbf{y}) = \mathbb{E}[\mathrm{var}[\widetilde{Z}_0|\widetilde{X}_0]|\widetilde{Y}_0 = \mathbf{y}]$. *Furthermore, if* $\sigma^2(m) > 0$, *then* A2(a) *holds with* $\sigma_3 = \sigma(m)$ *and, as* $n \to \infty$,

$$n^{1/2}(k_n(\mathbf{y}) - k(\mathbf{y})) \Rightarrow \frac{v(m)}{p(\mathbf{y})} N(0, 1),$$

where

$$v^2(m) = \sigma^2(m) + \frac{m-1}{m^2}\eta^2(\mathbf{y})p(\mathbf{y}).$$

Remark 4. As in Corollary 1, we can establish B3(b) by requiring that $(\widetilde{X}_i : i = 0, 1, \dots)$ is geometrically ergodic.

Remark 5. It should be noted that

$$\begin{aligned}\sigma^2(m) \to &\mathrm{var}[(\mu(X_0) - k(\mathbf{y}))I(Y_0 = \mathbf{y})] \\ &+ 2\sum_{i=1}^{\infty} \mathrm{cov}[(\mu(X_0) - k(\mathbf{y})) \\ &\times I(Y_0 = \mathbf{y}), (\mu(X_i) - k(\mathbf{y}))I(Y_i = \mathbf{y})]\end{aligned}$$

as $n \to \infty$.

4.1.1. How Do We Choose the Split Factor m?

We expect that the asymptotic variance of $k_n(\mathbf{y})$ will usually decrease as a function of the "split factor" m. Of course, the larger the value of m, the more computation time is spent. Consequently, we seek an optimal value $m^* \in \{1, 2, \dots\}$ that "trades off" the variance versus computational time. To compute m^*, we take the view that the computation time needed to generate X per unit time simulated is d_1 and the computation time needed to generate a single Z_{ij} is d_2. However, the Z_{ij}'s are generated only if $Y_i = \mathbf{y}$, so that the computation time c required to compute $k_n(\mathbf{y})$ is approximately $(d_1 + d_2 m)n$ (with $d_2 = d_2 p(\mathbf{y})$), so that $n \sim \lfloor c/(d_1 + d_2 m) \rfloor$ as $c \to \infty$. If $k_c(\mathbf{y})$ is the estimator available after expending c units of computer time, then we note that Theorem 3 implies that

$$c^{1/2}(k_c(\mathbf{y}) - k(\mathbf{y})) \Rightarrow \sqrt{(d_1 + d_2 m)v^2(m)/p(\mathbf{y})^2}\, N(0, 1) \quad (4.5)$$

as $c \to \infty$. We therefore choose m to minimize the asymptotic variance $h(m) \triangleq (d_1 + d_2 m)v^2(m)/p(\mathbf{y})^2$ appearing in (4.5). Because $h(\cdot)$ is eventually increasing, this minimization can be done by evaluating $h(\cdot)$ on some suitably chosen "interval" $\{1, 2, \dots, m'\}$ and choosing m^* as the minimizer of $h(\cdot)$ on this interval.

4.2. Continuous Case: Proposed Estimator $k_n'(\mathbf{y})$

We next turn to the setting in which $\widetilde{Y}_0$ is a continuous $\mathbb{R}^d$-valued rv with density $g(\cdot)$ for which $g(\mathbf{y}) > 0$. For the same choice of bandwidth $(a_n : n \geq 1)$ as in Section 3, let $\lambda_{in}(\mathbf{y}) = \lambda((Y_i - \mathbf{y})/a_n)$.

After generating a single run $((X_i, Y_i, Z_i) : i = 0, 1, \dots)$ along with the Z_{ij}'s, we estimate $k(\mathbf{y})$ via

$$k_n'(\mathbf{y}) = \frac{\sum_{i=0}^{n-1}\left(\frac{1}{m}Z_i + \frac{1}{m}\sum_{j=1}^{m-1} Z_{ij}\right)\lambda_{in}(\mathbf{y})}{\sum_{i=0}^{n-1} \lambda_{in}(\mathbf{y})}.$$

To establish the CLT for $k_n'(\mathbf{y})$, we use the following condition.

B4. B2 *holds with* $\widetilde{Z}_i$ *replaced by* $\overline{Z}_i \triangleq \frac{1}{m}Z_i + \frac{1}{m}\sum_{j=1}^{m-1} Z_{ij}$, $\widetilde{Y}_i$ *replaced by* Y_i, *and* $\widetilde{X}_i$ *replaced by* $(X_i, \overline{Z}_i)$ *for* $i = 0, 1, \dots$.

With B4 in place, we have a CLT for $k_n'(\mathbf{y})$. The proof of Theorem 4 is provided in Section EC.1 of the e-companion to this paper.

Theorem 4. *Let* $(X_i : i = 0, 1, \dots)$ *be a Markov chain with the same transition kernel as that of* $(\widetilde{X}_i : i = 0, 1, \dots)$. *Suppose that* $(X_i : i = 0, 1, \dots)$ *is initialized from its stationary distribution. Let* $(Z_{ij} : i = 0, 1, \dots, j = 1, 2, \dots)$ *be a family of real-valued rv's satisfying* (4.1).

Under A0 *and* B4, *as* $n \to \infty$,

$$(na_n^d)^{1/2}(k_n'(\mathbf{y}) - k(\mathbf{y})) \Rightarrow v'(m)N(0, 1),$$

where

$$\begin{aligned}v'^2(m) = &\left(\mathbb{E}[(\mu(\widetilde{X}_i) - k(\mathbf{y}))^2 \mid \widetilde{Y}_i = \mathbf{y}]\right. \\ &\left. + \frac{1}{m}\mathbb{E}[(\widetilde{Z}_i - \mu(\widetilde{X}_i))^2 \mid \widetilde{Y}_i = \mathbf{y}]\right)\frac{\int_0^1 \{f_\lambda(x)\}^2 dx}{g(\mathbf{y})\left(\int_0^1 f_\lambda(x)dx\right)^2}.\end{aligned}$$

Remark 6. When $a_n \sim an^{-\frac{1}{d+4}}$ for some positive real number a, the rate of convergence of $k_n'(\mathbf{y})$ is of order $n^{-\frac{2}{d+4}}$. Hence, the convergence rate degenerates as the dimension of S' increases.

Remark 7. The numerical behavior of $k_n'(\mathbf{y})$ depends on what type of kernel λ is used. In Section EC.7 of the e-companion to this paper, we observe the numerical behavior of $k_n'(\mathbf{y})$ for different types of kernels.

4.2.1. How Do We Choose the Split Factor m When $a_n \sim an^{-\gamma}$ for Some $0 < \gamma < 1/d$ as $n \to \infty$? To illustrate how the split factor m is chosen in practice, we consider a specific setting where $a_n \sim an^{-\gamma}$ as $n \to \infty$ for some positive real numbers a and γ satisfying $\gamma < 1/d$.

If $k_c'(\mathbf{y})$ is the version of $k_n'(\mathbf{y})$ available after expending c units of computer time (with the bandwidth for computer budget c chosen so that $a_n \sim an^{-\gamma}$ with $n = \lfloor c/(d_1 + d_2 m) \rfloor$), then we note that

$$c^{(1-\gamma d)/2}(k_c'(\mathbf{y}) - k(\mathbf{y})) \Rightarrow (d_1 + d_2 m)^{(1-\gamma d)/2} \frac{\nu'(m)}{a^{d/2}} N(0,1)$$

as $c \to \infty$. We therefore choose m to minimize the asymptotic mean square error $h'(m) \triangleq (d_1 + d_2 m)^{1-\gamma d} \nu'^2(m)/a^d$. Because $h'(\cdot)$ is eventually increasing, this minimization can be done by evaluating $h'(\cdot)$ on some suitably chosen interval $\{1, 2, \ldots, m'\}$ and choosing m^* as the minimizer of $h'(\cdot)$ on this interval.

4.2.2. An Alternative Estimator. We conclude this section with a discussion of an alternative estimator to $k_n'(\mathbf{y})$. Suppose that we wish to sample the Z_{ij}'s in such a way that more Z_{ij} samples are taken whenever Y_i is close to $\mathbf{y}$ (say, in proportion to $\lambda((\mathbf{y} - Y_i)/a_n)$. In particular, let $N_1, N_2, \ldots$ be a sequence of independent and identically distributed (i.i.d.) unit rate Poisson processes independent of $((X_i, Y_i, Z_i) : i = 0, 1, \ldots)$ and the Z_{ij}'s. Given $((X_i, Y_i, Z_i) : i = 0, 1, \ldots)$, for $\theta > 0$, $N_i(\theta \lambda_{in}(\mathbf{y}))$ follows a Poisson distribution with a mean of $\theta \lambda_{in}(\mathbf{y})$. For $\theta > 0$, we suggest the following estimator

$$\bar{k}_n'(\mathbf{y}) = \frac{\sum_{i=0}^{n-1} \sum_{j=1}^{N_i(\theta \lambda_{in}(\mathbf{y}))} Z_{ij}}{\sum_{i=0}^{n-1} N_i(\theta \lambda_{in}(\mathbf{y}))}$$

as an alternative to $k_n'(\mathbf{y})$.

The CLT for $\bar{k}_n'(\mathbf{y})$, which is provided in Section EC.2 of the e-companion to this paper, states that $\bar{k}_n'(\mathbf{y})$ asymptotically behaves identically to $\sum_{i=0}^{n-1}(\mu(X_i) - k(\mathbf{y}))\lambda_{in}(\mathbf{y}) / \sum_{i=0}^{n-1} \lambda_{in}(\mathbf{y})$.

The numerical performance of $\bar{k}_n'(\mathbf{y})$ is compared with that of $k_n'(\mathbf{y})$ in EC.8 of the e-companion to this paper.

4.2.3. Selecting the Bandwidth Sequence $(a_n : n \geq 1)$. The performance of $k_n'(\mathbf{y})$ is highly sensitive to the choice of the bandwidth sequence $(a_n : n \geq 1)$. One way to select the bandwidth is via leave-one-out cross validation. We introduce two heuristic approaches that are based on leave-one-out cross validation. In the first, after generating (X_i, Y_i, Z_i) along with the Z_{ij}'s for $i = 0, 1, \ldots, n-1$ and $j = 1, \ldots, m-1$, we treat $((Y_i, \bar{Z}_i) : i = 0, 1, \ldots, n-1)$ as a data set used for estimating the unknown function $k(\cdot)$, where $\bar{Z}_i = \frac{1}{m} Z_i + \frac{1}{m} \sum_{j=1}^{m-1} Z_{ij}$ for $i = 0, 1, \ldots, n-1$. Using the idea of leave-one-out cross validation, we select a_n as the minimizer of

$$\mathrm{CV}(\zeta) \triangleq \frac{1}{n} \sum_{i=0}^{n-1} (k_{-i,\zeta}'(Y_i) - \bar{Z}_i)^2$$

over $\zeta > 0$, where $k_{-i,\zeta}'(Y_i)$ is calculated from

$$k_{-i,\zeta}'(Y_i) = \frac{\sum_{k=0, k \neq i}^{n-1} \bar{Z}_k \lambda((Y_k - Y_i)/\zeta)}{\sum_{k=0, k \neq i}^{n-1} \lambda((Y_k - Y_i)/\zeta)}.$$

In other words, $k_{-i,\zeta}'(Y_i)$ is our proposed estimator of $k(Y_i)$ calculated from $((Y_k, \bar{Z}_k) : k = 0, 1, \ldots, n-1, k \neq i)$ using the bandwidth ζ. It should be noted that the Z_{ij}'s should be generated for every $i = 0, 1, \ldots, n-1$ in this approach, so it may require a significant amount of time to compute $\mathrm{CV}(\zeta)$. One way to reduce the computation time is by selecting a_n via cross validation using $((Y_i, Z_i) : i = 0, 1, \ldots, n-1)$ only. Thus, in the second approach, we let a_n be the minimizer of

$$\mathrm{CV}_*(\zeta) \triangleq \frac{1}{n} \sum_{i=0}^{n-1} (k_{-i,\zeta,*}'(Y_i) - Z_i)^2$$

over $\zeta > 0$, where $k_{-i,\zeta,*}'(Y_i)$ is calculated from

$$k_{-i,\zeta,*}'(Y_i) = \frac{\sum_{k=0, k \neq i}^{n-1} Z_k \lambda((Y_k - Y_i)/\zeta)}{\sum_{k=0, k \neq i}^{n-1} \lambda((Y_k - Y_i)/\zeta)}.$$

In Section EC.9 of the e-companion to this paper, we report the numerical behavior of $k_n'(\mathbf{y})$ when $(a_n : n \geq 1)$ is selected using these two approaches. For a general discussion on how the bandwidth is selected via cross validation in the kernel regression setting, see section 8 of Györfi et al. (2002).

Remark 8. In choosing the bandwidth sequence $(a_n : n \geq 1)$, one possible goal is to minimize the asymptotic mean square error. In this paper, we choose the bandwidth so as to make the bias term asymptotically negligible relative to the variance. We do this because our goal here is to produce confidence intervals for our prediction based on limiting normal rv's with zero mean. To produce estimators with minimal mean square error, the bandwidth must be chosen so that the limiting normal includes a nonzero bias term (see theorem 5 of Masry 2005).

5. Extensions

5.1 Markov Processes with Large Discrete State Space: Proposed Estimator $k_n^*(\mathbf{y})$

When S' is discrete, but large, it may be that $p(\mathbf{y})$ is still positive but very small. In this setting, the estimator $k_n(\mathbf{y})$ may be unusable from a practical point of view, unless $n \gg 1/p(\mathbf{y})$, since $np(\mathbf{y})$ is the expected

number of visits to $\mathbf{y}$ over n time units. In this setting, we propose using Z_i even if Y_i is not exactly equal to $\mathbf{y}$. In particular, we can assign a weight of $\lambda((\mathbf{y}-Y_i)/a_n)$ to the instance of Z_i, so more weight is applied to Z_i when Y_i is closer to $\mathbf{y}$. This leads to the following new estimator of $k(\mathbf{y})$ for the case when S' is discrete:

$$k_n^*(\mathbf{y}) = \frac{\sum_{i=0}^{n-1}\left(\frac{1}{m}Z_i + \frac{1}{m}\sum_{j=1}^{m-1} Z_{ij}\right)\lambda_{in}(\mathbf{y})}{\sum_{i=0}^{n-1}\lambda_{in}(\mathbf{y})}.$$

The CLT for $k_n^*(\mathbf{y})$, which is provided in Section EC.3 of the e-companion to this paper, states that $k_n^*(\mathbf{y})$ asymptotically behaves identical to $k_n(\mathbf{y})$. A more detailed rationale behind the development of $k_n^*(\mathbf{y})$ is provided in Section EC.4 of the e-companion to this paper.

5.2. Confidence Interval Methodology

We now briefly describe a confidence interval methodology for the estimators described in Sections 3 and 4. We illustrate this in the setting of the estimator $k_n(\mathbf{y})$, but the same approach can be implemented for any of our estimators.

The approach that we will follow is to use the method of batch means. In particular, we view $k_n(\mathbf{y})$ as a functional of $((X_i, Z_{ij}) : i = 0,1,\ldots,n-1, j = 1,2,\ldots m-1)$, namely, $k_n(\mathbf{y}) = \rho(X_i, Z_{ij} : i = 0,1,\ldots,n-1, j = 1,2,\ldots m-1)$ for some appropriately defined mapping $\rho(\cdot)$. Under the assumptions made earlier in this paper, we expect X to "mix" rapidly enough that the $k_{i,n}(\mathbf{y})$'s are asymptotically independent, where

$$k_{i,n}(\mathbf{y}) = \rho(X_{(i-1)n+v}, Z_{(i-1)n+v,j} : v = 0,1,\ldots,n-1, \quad j = 1,2,\ldots,m-1)$$

for $i = 1,2,\ldots$ is obtained by forming the estimator defined by ρ, based on the observations associated with the ith batch having length n.

Theorem 3 states that

$$n^{1/2}(k_n(\mathbf{y}) - k(\mathbf{y})) \Rightarrow \nu(m)/p(\mathbf{y})N(0,1)$$

as $n \to \infty$. In the presence of appropriate mixing assumptions, we expect that

$$n^{1/2}(k_{1,n}(\mathbf{y}) - k(\mathbf{y}), \ldots, k_{l,n}(\mathbf{y}) - k(\mathbf{y})) \Rightarrow \nu(m)/p(\mathbf{y})(N_1(0,1), \ldots, N_l(0,1))$$

as $n \to \infty$, where the $N_i(0,1)$ rv's are independent normal rv's with mean 0 and unit variance. As a consequence, the continuous mapping principle implies that, for $l \geq 2$,

$$\frac{l^{1/2}\left(\frac{1}{l}\sum_{i=1}^{l} k_{i,n}(\mathbf{y}) - k(\mathbf{y})\right)}{\sqrt{\frac{1}{l-1}\sum_{i=1}^{l}\left(k_{i,n}(\mathbf{y}) - \frac{1}{l}\sum_{j=1}^{l} k_{j,n}(\mathbf{y})\right)^2}} \Rightarrow t_{l-1}$$

as $n \to \infty$, where t_{l-1} is a student's t-distribution rv with $l-1$ degrees of freedom. Hence, the interval

$$\left[\frac{1}{l}\sum_{i=1}^{l} k_{i,n}(\mathbf{y}) - \frac{zs_n}{\sqrt{l}}, \frac{1}{l}\sum_{i=1}^{l} k_{i,n}(\mathbf{y}) + \frac{zs_n}{\sqrt{l}}\right]$$

is an asymptotic $100(1-\delta)\%$ confidence interval for $k(\mathbf{y})$ (with asymptotically exact converge as $n \to \infty$), where

$$s_n = \sqrt{\frac{1}{l-1}\sum_{i=1}^{l}\left(k_{i,n}(\mathbf{y}) - \frac{1}{l}\sum_{j=1}^{l} k_{j,n}(\mathbf{y})\right)^2}$$

and z is chosen so that $\mathbb{P}(-z \leq t_{l-1} \leq z) = 1-\delta$. Similar batch means confidence intervals, based on the student's t-distribution rv with $l-1$ degrees of freedom, can be constructed for all of the other estimators proposed in this paper.

5.3. Estimating $\mathbb{E}[\widetilde{Z}(s+t) \mid \widetilde{Y}(s) = \mathbf{y}]$ for Various Values of y and t

It should be noted that the proposed methods can be extended to the case where one wishes to estimate $\mathbb{E}[\widetilde{Z}(s+t)|\widetilde{Y}(s) = \mathbf{y}]$ for various values of $\mathbf{y}$ and t. Specifically, once a long simulation run $((X_i, Y_i, Z_i) : i = 0,1,\ldots)$ is generated, this simulation run can be used to compute $\mathbb{E}[\widetilde{Z}(s+t)|\widetilde{Y}(s) = \mathbf{y}]$ for various values of $\mathbf{y}$ and t. When we want to compute $\mathbb{E}[\widetilde{Z}(s+t)|\widetilde{Y}(s) = \mathbf{y}]$ for various values of $\mathbf{y}$, only the Z_{ij}'s need to be generated again for each $\mathbf{y}$. Since the Z_{ij}'s need to be generated only when $Y_i = \mathbf{y}$ or $\lambda((Y_i - \mathbf{y})/a_n)$ is positive, the amount of time required to generate the Z_{ij}'s is small compared with the amount of time required to generate $((X_i, Y_i, Z_i) : i = 0,1,\ldots)$. Computing $\mathbb{E}[\widetilde{Z}(s+t)|\widetilde{Y}(s) = \mathbf{y}]$ for various values of t also requires generating the Z_{ij}'s only again.

6. Computational Examples

We conducted all simulations in this section using a 64-bit computer with an Intel Core i7-6700K CPU at 4 GHz and a RAM of 32 GB, and programmed all simulations in MATLAB R2018a.

6.1. Continuous Case

We first consider the case where $\widetilde{Y}(s)$ is a continuous rv. Specifically, we consider a stationary two-dimensional Ornstein–Uhlenbeck process in which $\widetilde{X} = ((\widetilde{X}_1(u), \widetilde{X}_2(u)) : u \geq 0)$ satisfies the pair of stochastic differential equations

$$d\widetilde{X}_1(u) = -\widetilde{X}_1(u)du + dB_1(u)$$
$$d\widetilde{X}_2(u) = -\widetilde{X}_1(u)du - \widetilde{X}_2(u)du + dB_2(u),$$

where $\widetilde{B}_1 = (\widetilde{B}_1(u) : u \geq 0)$ and $\widetilde{B}_2 = (\widetilde{B}_2(u) : u \geq 0)$ are two independent standard Brownian motions.

Our goal is to estimate $k(2) = \mathbb{E}[\widetilde{Z}(s+1) \mid \widetilde{Y}(s) = 2]$, where $\widetilde{Z}(u) = \widetilde{X}_2(u)$ and $\widetilde{Y}(u) = \widetilde{X}_1(u)$ for $u \geq 0$. In this case, $k(2)$ can be computed analytically and is given by $k(2) = -1.10$ (see Section EC.6 of the e-companion to this paper for the detailed derivation of $k(2) = -1.10$).

6.1.1. Confidence Intervals Based on the Method of Batch Means. For fixed n, we used $a_n = 0.5n^{-1/5}$ and $\lambda(z) = 1/2$ if $-1 \leq z \leq 1$, and 0 otherwise. We next generated (X, Y, Z) along with the Z_{ij}'s with X initialized at (0, 0). We computed the 95% confidence intervals for $k(2)$ using $k_n'(2)$ and the method of batch means described in Section 5.2 with $l = 20$ batches. We replicated the confidence interval procedure 200 independent times. Table 1 reports the average of the 200 computer times required to generate each confidence interval measured in seconds (T), the average of the 200 confidence interval midpoints, the average of the 200 half-widths, and the proportion (P) of time that the 200 confidence intervals covered the true value of $k(2)$, when $m = 1, 4,$ and 8. In Table 1, CI is the average confidence interval midpoint $\pm$ the average confidence interval half-width. Table 1 reveals that the proportion of the 95% confidence intervals covering the true value of $k(2)$ converges to its target value, which is 95%, as n increases.

6.1.2. Effect of the Split Factor *m*. Theorem 4 describes how the split factor m affects the asymptotic mean square error $v'(m)^2$. Specifically, Theorem 4 states that $v'(m)^2$ decreases as m increases. To illustrate this numerically, we generated a copy of $k_n'(2)$ using $a_n = 0.5n^{-1/5}$ and $\lambda(z) = 1/2$ if $-1 \leq z \leq 1$ and 0 otherwise, with X initialized at (0, 0). We then computed $(k_n'(2) - k(2))^2$ as an estimate of the mean square error (MSE) of $k_n'(2)$. We replicated this procedure 200 independent times, generating 200 independent estimates of the MSE. We then computed the 95% confidence interval of the MSE using these 200 independent values. Table 2 reports these 95% confidence intervals (CI) for a variety of the m values when n = 1,000 and n = 2,000, respectively. Table 2 also reports the computer time (T) required to generate a single copy of $k_n'(2)$ for each m and n value. Table 2 shows that, as m increases, the MSE of $k_n'(2)$ decreases, as Theorem 4 claimed. It should be noted that more computer time is required to generate $k_n'(2)$ as m increases. Thus, one can expect that there is an optimal value of m minimizing the asymptotic mean square error of $k_n'(2)$ given a fixed computational budget.

6.1.3. How Do We Verify the Convergence Rate Empirically? Theorem 4 establishes that the convergence rate of $(k_n'(2) - k(2))^2$ is of order $n^{-4/5}$ when $a_n = O(n^{-1/5})$. To verify this empirically, we note that $(k_n'(2) - k(2))^2 = O(n^{-4/5})$ is verified if $\log(k_n'(2) - k(2))^2 = -(4/5)\log(n) + O(1)$ as $n \to \infty$ is verified. Thus, the problem boils down to the question of verifying that the slope of the graph of $\log(k_n'(2) - k(2))^2$ against $\log(n)$ is $-4/5$. Toward this goal, for each n, we generated a copy of $k_n'(2)$ using $m = 1$, $a_n = 0.5n^{-1/5}$, and the uniform density $\lambda(\cdot)$ over $[-1, 1]$; computed $(k_n'(2) - k(2))^2$; and repeated this procedure 20 independent times, generating 20 independent copies of $(k_n'(2) - k(2))^2$. We then computed the average of these 20 copies, which is denoted by $\overline{M}_n$. Table 3 shows the $\overline{M}_n$ values for $n \in \{500, 1,000, 5,000, 10,000, 50,000, 100,000\}$.

We plotted the logarithm of $\overline{M}_n$ as the dependent variable, against the logarithm of the n values as the independent variable, and computed the slope of this graph using the least squares method. The slope, computed this way, is –0.799, which is very close to its theoretical value –0.8.

6.1.4. How Does the Kernel-Based Estimator $k_n'(\mathbf{y})$ Compare with Global Surface Fitting? In this example, S' is a continuous space. Thus, it is possible to construct a "response surface" $\widetilde{k}_n(\cdot)$ of $k(\cdot)$ and use $\widetilde{k}_n(\mathbf{y})$ as an estimator of $k(\mathbf{y})$. (This cannot be easily generalized to the discrete case.) Once one generates (X_i, Y_i, Z_i) for $i = 0, 1, \ldots$ along with the Z_{ij}'s, such a response surface can be constructed by using the Y_i's as the data for the independent variable and $\overline{Z}_i \triangleq Z_i/m + \sum_{j=1}^{m-1} Z_{ij}/m$ for $i = 1, 2, \ldots$ as the data for the dependent variable. A key decision here is which metamodel is used for the response surface of $k(\cdot)$. In this example, we expect that $k(\cdot)$ is smooth enough so that $k(\cdot)$ has a square-integrable second derivative, namely, $\int_{-\infty}^{\infty} \{k^{(2)}(\mathbf{z})\}^2 d\mathbf{z} < \infty$. Thus, as a response surface of $k(\cdot)$, we use the solution to the following minimization problem,

$$\underset{f \in \mathcal{F}}{\text{minimize}} \frac{1}{n} \sum_{i=0}^{n-1} (\overline{Z}_i - f(Y_i))^2 + c \int_{-\infty}^{\infty} \{f^{(2)}(\mathbf{z})\}^2 d\mathbf{z} \qquad (6.1)$$

for some smoothing constant $c > 0$, where $\mathcal{F}$ is the set of functions $f : \mathbb{R} \to \mathbb{R}$, whose first derivative is

Table 1. 95% Confidence Intervals of $k_n'(2)$ Based on the Batch Means Method

	$m = 1$			$m = 4$			$m = 8$		
n	T	CI	P	T	CI	P	T	CI	P
1,000	1.5	-1.00 ± 0.25	0.85	1.5	-1.02 ± 0.17	0.88	1.5	-1.00 ± 0.17	0.85
2,000	3.1	-1.08 ± 0.28	0.95	3.1	-1.08 ± 0.11	0.92	3.0	-1.08 ± 0.10	0.91
4,000	6.2	-1.10 ± 0.13	0.95	6.2	-1.09 ± 0.08	0.94	6.0	-1.10 ± 0.06	0.95
True value		−1.10			−1.10			−1.10	

Table 2. 95% Confidence Intervals of the Mean Square Error of $k_n'(2)$

	n = 1,000		n = 2,000	
m	T	CI	T	CI
1	0.08	0.31 ± 0.06	0.16	0.22 ± 0.06
4	0.08	0.18 ± 0.05	0.16	0.07 ± 0.02
40	0.09	0.16 ± 0.05	0.17	0.04 ± 0.02
400	0.16	0.14 ± 0.05	0.30	0.03 ± 0.02
4,000	1.07	0.11 ± 0.04	1.68	0.03 ± 0.01

absolutely continuous and the second derivative $f^{(2)}$ satisfies $\int_{-\infty}^{\infty}\{f^{(2)}(\mathbf{z})\}^2 d\mathbf{z} < \infty$. The solution $\widetilde{k}_n(\cdot)$ to (6.1) is referred to as the *cubic smoothing spline*. It is well known that $\widetilde{k}_n(\cdot)$ can be expressed as a linear combination of some basis functions and that the coefficients in the linear combination can be found by solving a system of linear equations (see, e.g., p. 31 of Wahba 1990). We next compare the numerical behavior of $\widetilde{k}_n(2)$ to that of our proposed estimator $k_n'(2)$. For each m and n, we generated (X, Y, Z) along with the Z_{ij}'s; X was initialized at (0, 0). We computed $k_n'(2)$ using $a_n = 0.5n^{-1/5}$ and the uniform density $\lambda(\cdot)$ over $[-1, 1]$. We also computed $\widetilde{k}_n(2)$ by solving (6.1) with the smoothing constant c = 0.01. We then computed $(k_n'(2) - k(2))^2$ and $(\widetilde{k}_n(2) - k(2))^2$ as estimates of the MSE of $k_n'(2)$ and $\widetilde{k}_n(2)$, respectively. We repeated this procedure 1,000 independent times, generating 1,000 independent copies of $(k_n'(2) - k(2))^2$ and $(\widetilde{k}_n(2) - k(2))^2$. Table 4 reports the 95% confidence interval of the MSE of $k_n'(2)$ and $\widetilde{k}_n(2)$, computed from these 1,000 independent copies.

Table 4 reveals that the MSE of $\widetilde{k}_n(2)$ decreases as n increases. This justifies the validity of $k_n(2)$ as an estimator of $k(2)$. Table 4 also displays that the MSE of $\widetilde{k}_n(2)$ decreases as m increases. Since more computer time is required as m increases, one can expect that there is an optimal value of m minimizing the MSE of $\widetilde{k}_n(2)$ for a fixed computational budget. In other words, the "splitting" idea works to reduce the MSE, even if one generates an estimator of $k(2)$ using global surface fitting. It should be also noted that, for every value of n and m in Table 4, $\widetilde{k}_n(2)$ produces a lower MSE than $k_n'(2)$ does. This phenomenon is intuitively acceptable since $\widetilde{k}_n(2)$ makes use of the fact that $k(\cdot)$ is a smooth function over its domain, whereas $k_n'(2)$ does not make any functional assumptions on $k(\cdot)$.

Table 3. Average of 20 Independent Copies of $(k_n'(2) - k(2))^2$

n	500	1,000	5,000	10,000	50,000	100,000
$\overline{M}_n$	0.512	0.248	0.071	0.046	0.013	0.007

6.2. Discrete Case: Container Management System

In this section, we consider a simplified version of the container management system (see Figure 2), which motivated this paper. In this case, S' is discrete, but large, so we can use $k_n(\mathbf{y})$ and $k_n^*(\mathbf{y})$ to estimate $k(\mathbf{y})$.

6.2.1. Description of the Container Management System. We consider a supply chain of a car manufacturer that consists of six external suppliers, three manufacturing plants, and two assembly plants. Outside the supply chain, there are two external customers, who are dealerships A and B. Manufacturing plant A manufactures the main engine part. It has two external suppliers, suppliers A and B. Suppliers A and B supply brakes and other engine components, respectively. Manufacturing plant B manufactures the body parts. It has two external suppliers, suppliers C and D. Suppliers C and D supply doors and windows, respectively. Manufacturing plant C manufactures the interiors. It has two external suppliers, suppliers E and F. Suppliers E and F supply floor mats and car seats, respectively. Assembly plants A and B receive all of the parts from manufacturing plants A, B, and C and assemble them into the final products.

Dealership A places orders with assembly plant A only. The interarrival times of the orders coming from dealership A follow a discrete uniform distribution over one and two days. Each order placed by dealership A consists of 10 vehicles. Dealership B places orders with assembly plant B only. Interarrival times of the orders coming from dealership B follow a discrete uniform distribution over one and two days. Each order placed by dealership B consists of 10 vehicles.

When an assembly plant receives an order of 10 vehicles, it immediately places an order of the 10 vehicle equivalents with manufacturing plants A, B, and C. In other words, if a vehicle requires two interiors, then the assembly plant places an order of 20 interiors with manufacturing plant C. Whenever this happens, each manufacturing plant places an order of the 10 vehicle equivalents with each of its external suppliers immediately.

We assume that each external supplier has 10 trucks, each manufacturing plant has 5 trucks, and each assembly plant has 4 trucks. Each truck can carry a container that can hold 10 vehicle equivalents.

We also assume that each external supplier has an unlimited quantity of parts available. When an external supplier receives an order, he or she checks if there is a truck available. If there is a truck available, then the external supplier loads the truck with the requested parts and ships them out to the destination plant. If there are not any trucks available, then the order is backordered. The time between when an external supplier receives an order and when he or she

Table 4. 95% Confidence Interval of the MSE of $k_n'(2)$ and $\widetilde{k}_n(2)$

	$m = 1$		$m = 2$		$m = 4$	
n	$k_n'(2)$	$\widetilde{k}_n(2)$	$k_n'(2)$	$\widetilde{k}_n(2)$	$k_n'(2)$	$\widetilde{k}_n(2)$
100	0.960 ± 0.035	0.204 ± 0.021	0.870 ± 0.033	0.190 ± 0.018	0.867 ± 0.032	0.176 ± 0.015
200	0.821 ± 0.040	0.097 ± 0.009	0.733 ± 0.034	0.092 ± 0.008	0.661 ± 0.035	0.091 ± 0.009
500	0.558 ± 0.039	0.040 ± 0.004	0.429 ± 0.032	0.039 ± 0.003	0.370 ± 0.031	0.034 ± 0.003

ships out the order is negligible, given that a truck is available. When a truck arrives at the destination plant, it unloads the parts immediately, loads the empty container back, and returns to the original plant. The time required for a truck to travel from an external supplier to a manufacturing plant (or from a manufacturing plant to an external supplier) follows a discrete uniform distribution between 5 and 25 days.

Each manufacturing plant waits until all of the parts for an order are received. Once all of the parts are received, they join a queue of orders, each of which will be served by a single server. Once each order is completed by the server, the manufacturing plant checks if there is a truck available. If a truck is available, then the order is loaded and shipped to the destination assembly plant. If there are not any trucks available, then the order waits until a truck becomes available. Each manufacturing plant can be modeled as a single-server queue with a buffer of limited capacity, where each job is served based on the first-in/first-out (FIFO) discipline. Each manufacturing plant keeps a list of where each order came from so that it can deliver the completed orders to the destination assembly plants on the list. If the list has 20 orders, then the manufacturing plant rejects any additional orders and the rejected orders are lost. Thus, the list always has at most 20 orders. The time required to process each order at the server of a manufacturing plant follows a discrete uniform distribution between two and eight days in manufacturing plants A, B, and C. The time required for a truck to travel from a manufacturing plant to an assembly plant (or from an assembly plant to a manufacturing plant) follows a discrete uniform distribution between 6 and 20 days.

Assembly plants work in a similar fashion. Each assembly plant waits until all of the parts for an order are received. Once all of the parts are received, they join a queue of orders, each of which will be served by a single server. Once an order is completed by the server, the assembly plant checks if there is a truck available. If a truck is available, then the order is loaded and shipped to the destination dealership. If there are not any trucks available, then the order waits until a truck becomes available. Each assembly plant can be modeled as a single-server queue with a buffer of infinite capacity, where each job is served based on the FIFO discipline. The time required to process each order at the server of each assembly plant follows a discrete uniform distribution between 1 and 18 days. The time required for a truck to travel from an assembly plant to a dealership (or from a dealership to an assembly plant) follows a discrete uniform distribution between 1 and 40 days.

6.2.2. Description of the Underlying Markov Process. The supply chain described in Section 6.2.1 can be simulated by recursively updating the state variables of a Markov process $\widetilde{X} = (\widetilde{X}_i : i = 0, 1, \dots)$, where $\widetilde{X}_i$ consists of the following state variables:

- the elapsed interarrival times of orders coming from dealerships A and B,
- the number of jobs being served or waiting in line at each manufacturing or assembly plant,
- the list of orders, at each manufacturing plant, that tells us where each order came from and the total number of orders on the list,
- the elapsed service time of the order that is currently being served at each manufacturing plant and each assembly plant,
- the numbers of containers moving from each external supplier to each manufacturing plant, from each manufacturing plant to each external supplier, from each manufacturing plant to each assembly plant, from each assembly plant to each manufacturing plant, from each assembly plant to each dealership, from each dealership to each assembly plant,
- the elapsed traveling time of each container in transit,
- the number of brakes in manufacturing plant A waiting for the other parts to arrive,
- the number of other engine components in manufacturing plant A waiting for the other parts to arrive,
- the number of doors in manufacturing plant B waiting for the other parts to arrive,
- the number of windows in manufacturing plant B waiting for the other parts to arrive,
- the number of floor mats in manufacturing plant C waiting for the other parts to arrive,
- the number of car seats in manufacturing plant C waiting for the other parts to arrive,
- the number of main engine parts in each assembly plant waiting for the other parts to arrive,
- the number of body parts in each assembly plant waiting for the other parts to arrive,

• the number of interiors in each assembly plant waiting for the other parts to arrive,
• the number of backorders at each external supplier,
• and the number of finished orders waiting for available containers at each manufacturing or assembly plant at the beginning of the ith day. We assume that we can observe $\widetilde{Y} = (\widetilde{Y}_i : i = 0, 1, \dots) = ((\widetilde{Y}_i^A, \widetilde{Y}_i^B, \widetilde{Y}_i^C) : i = 0, 1, \dots)$, where $\widetilde{Y}_i^A$, $\widetilde{Y}_i^B$, and $\widetilde{Y}_i^C$ are the total numbers of orders on the list of orders at manufacturing plants A, B, and C, respectively, at the beginning of the ith day. We also assume that $\widetilde{Z}_i$ is the number of completed orders waiting for available containers at manufacturing plant C at the beginning of the ith day. In this setting, we wish to estimate $\mathbb{E}(\widetilde{Z}_{15} \mid \widetilde{Y}_0 = \mathbf{y})$, where

$$\mathbf{y} = (20, 13, 14).$$

6.2.3. Numerical Behavior of $k_n(\mathbf{y})$ and $k_n^*(\mathbf{y})$. For fixed n, we used $a_n = 2n^{-1/7}$, suggested by Remark 6 analogously to the continuous case. It should be noted that Remark 6 suggests that $a_n \sim an^{-\frac{1}{d+4}}$, where d is the dimension of $\mathbf{y}$, not the dimension of the state space of $\widetilde{X}$; $\lambda(\mathbf{z})$ is the density function of the multivariate standard normal variable evaluated at $\mathbf{z} \in \mathbb{R}^3$ if $\|\mathbf{z}\| \le 5$, and 0 otherwise. We next generated (X, Y, Z) along with the Z_{ij}'s, with X initialized empty and idle. (The elapsed interarrival times of orders coming from dealerships A and B are set to be equal to 1.) We then computed $k_n(\mathbf{y})$ and $k_n^*(\mathbf{y})$, along with their MSE $(k_n(\mathbf{y}) - k(\mathbf{y}))^2$ and $(k_n^*(\mathbf{y})) - k(\mathbf{y}))^2$ and replicated this procedure 200 independent times. Table 5 reports the average (MSE) of 200 copies of $(k_n(\mathbf{y}) - k(\mathbf{y}))^2$ and $(k_n^*(\mathbf{y})) - k(\mathbf{y}))^2$, respectively. Table 6 reports the average amount of computer time (T) required to generate each copy of $k_n(\mathbf{y})$ (and $k_n^*(\mathbf{y})$, respectively) measured in seconds, the average (Mean) of the 200 copies of $k_n(\mathbf{y})$ (and $k_n^*(\mathbf{y})$, respectively), and the standard deviation (Std) of the 200 copies of $k_n(\mathbf{y})$ (and $k_n^*(\mathbf{y})$, respectively). The true value in the last row of Table 6 is estimated by averaging 20 i.i.d. copies of $k_n(\mathbf{y})$ with $m = 40$ and $n = 500{,}000$.

Tables 5 and 6 reveal that, as m increases, more computer time is required to generate each estimator, but a lower mean square error is achieved. They also suggest that $k_n^*(\mathbf{y})$ outperforms $k_n(\mathbf{y})$, especially when n is relatively small.

Table 5. MSE of $k_n(\mathbf{y})$ and $k_n^*(\mathbf{y})$ in the Container Management Example When $m = 1$ and $m = 4$

	n	$k_n(\mathbf{y})$	$k_n^*(\mathbf{y})$
$m = 1$	20,000	13.8	9.5
	40,000	6.6	4.4
	80,000	2.7	2.1
$m = 4$	20,000	13.3	8.4
	40,000	5.8	3.5
	80,000	1.7	1.2

Table 6. Computer Time, Average, and Standard Deviation of $k_n(\mathbf{y})$ and $k_n^*(\mathbf{y})$ in the Container Management Example When $m = 1$ and $m = 4$

	$m = 1$						$m = 4$					
	$k_n(\mathbf{y})$			$k_n^*(\mathbf{y})$			$k_n(\mathbf{y})$			$k_n^*(\mathbf{y})$		
n	T	Mean	Std	T	Mean	Std	T	Mean	Std	T	Mean	Std
20,000	30	4.8	3.5	30	6.2	3.1	36	4.4	3.3	206	5.7	2.9
40,000	68	6.3	2.6	68	6.8	2.0	78	5.7	2.4	242	6.2	1.9
80,000	174	6.9	1.4	174	6.9	1.1	192	6.3	1.3	455	6.3	1.0
True value		6.0			6.0			6.0			6.0	

7. Conclusions

Motivated by a container management system, we considered the problem of estimating $k(\mathbf{y}) = \mathbb{E}[\widetilde{Z}(s+t) \mid \widetilde{Y}(s) = \mathbf{y}]$, where $(\widetilde{X}(u) : u \ge 0)$ is the underlying Markov process, $(\widetilde{Y}(u) : u \ge 0)$ is the observed process, $\mathbf{y}$ is the observation made at the current time s, and $\widetilde{Z}(s+t)$ is what we want to predict for time $s + t$. The challenge was that the state space of $\widetilde{Y}(s)$ is discrete and large, and $\mathbf{y}$ is visited rarely. We tackled the issue of initialization, proposed estimators based on the idea of splitting and the kernel regression estimation, and established CLT-type results for the proposed estimators.

Offline-simulation online-application (OSOA) has the advantage that when decisions need to be made quickly in real time, the OSOA approach has the ability to quickly return predictions, because computing the prediction at a newly acquired $\mathbf{y}$ value only involves a function evaluation of the response surface. On the other hand, when sufficient computing time is available, our methodology provides fully rigorous large-sample confidence intervals for the prediction. Both methods suffer dimensional degradation when the dimension of $\mathbf{y}$ is large. Further research will be required to better understand the relative advantages and disadvantages of these newly introduced methodologies, especially in high dimensional settings.

Figure 2. A Simplified Version of the Supply Chain of a Car Manufacturer

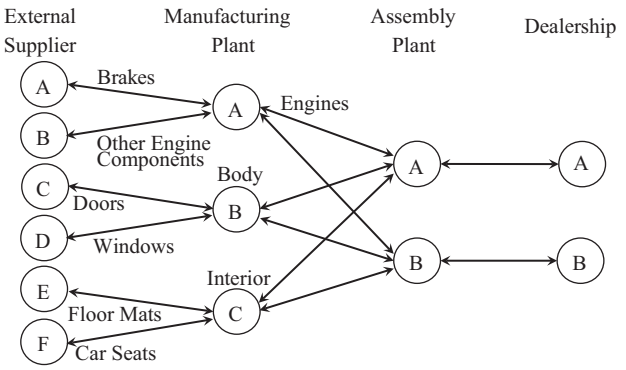

Acknowledgments

The authors express their gratitude to Shang-Tae Yee and Armin Zimmermann for technical discussions that led to this research topic. The authors also thank the referees and the associate editor for their careful reading of the paper and the suggestions they made that improved the paper. In addition, P. W. Glynn thanks the Hong Kong Institute for Advanced Study for its support and hospitality in conducting a part of this research.

References

Abdellah AB, L'Ecuyer P, Owen A, Puchhammer F (2018) Density estimation by randomized quasi–Monte Carlo. Preprint, submitted July 6, https://www.iro.umontreal.ca/~lecuyer/myftp/papers/density-rqmc-may2020.pdf.

Asmussen S (2003) *Applied Probability and Queues* (Springer, New York).

Asmussen S, Glynn PW (2007) *Stochastic Simulation: Algorithms and Analysis* (Springer, New York).

Bierens H (1987) Kernel estimators of regression functions. Bewley T, ed., *Adv. Econometrics Fifth World Congress* (Cambridge University Press, Cambridge, UK), 99–144.

Doob JL (1953) *Stochastic Processes* (Wiley, New York).

Doucet A, Johansen AM (2009) A tutorial on particle filtering and smoothing: Fifteen years later. Crisan D, Rozovskii B, eds. *Oxford Handbook of Nonlinear Filtering* (Oxford University Press, New York), 635–655.

Györfi L, Kohler M, Krzyżak A, Walk H (2002) *A Distribution-Free Theory of Nonparametric Regression* (Springer, New York).

Hong LJ, Jiang G (2019) Offline simulation online application: A new framework of simulation-based decision making. *Asia-Pacific J. Oper. Res.* 36(6):1940015.

Karlin S, Taylor HM (1975) *A First Course in Stochastic Processes* (Academic Press, New York).

Liu J, West M (2001) Combined parameter and state estimation in simulation-based filtering. Doucet A, de Freitas N, Gordon N, eds. *Sequential Monte Carlo Methods in Practice: Statistics for Engineering and Information Science* (Springer, New York), 197–223.

Masry E (2005) Nonparametric regression estimation for dependent functional data: Asymptotic normality. *Stochastic Processes Appl.* 115(1):155–177.

Nelson BL (2016) Some tactical problems in digital simulation' for the next 10 years. *J. Simulation* 10(1):2–11.

Nummelin E, Tuominen P (1982) Geometric ergodicity of Harris recurrent Markov chains with applications to renewal theory. *Stochastic Processes Appl.* 12(2):187–202.

Peng Y, Fu MC, Heidergott B, Lam H (2020) Maximum likelihood estimation by Monte Carlo simulation: Toward data-driven stochastic modeling. *Oper. Res.* 68(6):1896–1912.

Pitt MK, Shephard N (1999) Filtering via simulation: Auxiliary particle filter. *J. Amer. Statist. Assoc.* 94(446):590–599.

Rosenblatt M (1971) *Markov Processes. Structure and Asymptotic Behavior* (Springer, New York).

Wahba G (1990) *Spline Models for Observational Data* (SIAM, Philadelphia).

Williams D (1991) *Probability with Martingales* (Cambridge University Press, Cambridge, UK).

Eunji Lim is an assistant professor in the Department of Decision Sciences and Marketing at Adelphi University. Her research interests include simulation-based optimization and nonparametric function estimation under shape constraints.

Peter W. Glynn is the Thomas Ford Professor in the Department of Management Science and Engineering at Stanford University. His research interests lie in simulation, computational probability, queueing theory, statistical inference for stochastic processes, and stochastic modeling.