

International Series in
Operations Research & Management Science

Barry L. Nelson

Foundations and Methods of Stochastic Simulation

A First Course



 Springer

International Series in Operations Research & Management Science

Volume 187

Series Editor

Frederick S. Hillier
Stanford University, CA, USA

Special Editorial Consultant

Camille C. Price
Stephen F. Austin State University, TX, USA

This book was recommended by Dr. Price

For further volumes:

<http://www.springer.com/series/6161>

Barry L. Nelson

Foundations and Methods of Stochastic Simulation

A First Course

 Springer

Barry L. Nelson
Department of Industrial Engineering
and Management Sciences
Northwestern University
Evanston, IL, USA

ISSN 0884-8289

ISBN 978-1-4614-6159-3

ISBN 978-1-4614-6160-9 (eBook)

DOI 10.1007/978-1-4614-6160-9

Springer New York Heidelberg Dordrecht London

Library of Congress Control Number: 2012953267

© Springer Science+Business Media New York 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

*To Rowena Adelaar, Bahar Biller,
Justin Boesel, Marne Cario, Xi Chen,
Ira Gerhardt, Lynne Goldsman, Jeff Hong,
Seong-Hee Kim, Hai Lan, Soonhui Lee,
Ming Liu, Frank Matejcik, J. O. Miller,
Jamie Wieland Mulligan, Jutta Pichitlamken,
Tony Sharon, Yu-Hui Tao,
Mustafa Tongarlak, Shing Chih Tsai,
Hong Wan, Wei Xie, Jie Xu, Wendy Xu,
Feng Yang, Wei-Ning Yang, Mingjian Yuan,
and any students who follow. You made it all
worthwhile.*

Preface

As is often the case, this book grew out of a course. The interesting part is that the course grew out of this story: I was helping one of our graduate students, Viji Krishnamurthy, whose research involved developing rules for the use of flexible workers in a repair and maintenance environment (Iravani and Krishnamurthy 2007). Using Markov chain analysis, Viji had derived some optimal strategies for simple systems and now wanted to test their robustness for more realistic problems. This is where simulation came in. Viji had taken the typical first simulation modeling course that used a commercial simulation product and an advanced course that focused only on design and analysis, but not model building. She (and I) spent hours trying to trick a commercial product into simulating the complex worker allocation rules that she wanted to test. Commercial simulation environments make modeling easy by including the sort of system features that users typically want. Unfortunately, making it easy to model typical features can make it difficult to represent something different, and research is always about something different.

Finally, in frustration, I handwrote three pages of pseudocode to simulate exactly what Viji wanted, handed her the notes, and told her to code it up in C (which fortunately she could do). She came back the next day excited: “Now I can test anything I want, and it runs in seconds!” From this experience, IEMS 435 Introduction to Stochastic Simulation—a required course for all of our Ph.D. students—was born. IEMS 435 is not just about modeling and programming, however. Viji also needed to run well-designed experiments on her model, experiments that could provide compelling evidence for or against her analytically derived rules. So, experiment design and analysis are also a part of the IEMS 435 course, which this book was written to support.

The objectives of the book are as follows:

- To prepare students who have never had a discrete-event, stochastic simulation course to build simulations in a lower-level programming language. I am convinced that if they can do this, they can easily pick up higher-level simulation modeling environments when they need them (maybe for teaching a course as faculty).

- To prepare students to *use* simulation in their non-simulation research. This is why I emphasize actually programming simulations—which provides the greatest flexibility, control, and understanding—and experiment design and analysis.
- To prepare students to go into an advanced course on simulation methodology, including independent studies directed by their advisers. The usual first course in simulation emphasizes modeling and commercial software and is poor preparation for a research-oriented advanced course which may treat the simulation almost entirely as a mathematical object, breaking the critical connection between modeling and analysis.
- To provide a solid mathematical/statistical grounding in simulation and some (but not all) tools to solve actual problems.

The philosophy of the book is similar to Law (2007) in covering both simulation modeling and analysis, but is different in that there is no attempt to be comprehensive or survey the field. The goal is to be *concise, precise, and integrated*, leaving a lot of room for the instructor to expand on areas of interest or importance to them. The hope is that an instructor will want students to read *all* of the book to get a complete, coherent picture before jumping off into other reference texts or journal papers. To that end, I provide pointers to relevant literature. However, while not comprehensive, the book is *complete*; so, it is appropriate for students or researchers who need to learn the basics of simulation on their own without the benefit of a course. The book by Asmussen and Glynn (2007) shares some of the same objectives as this book; it is an excellent introduction to advanced simulation analysis and covers more of the topic than I do, but it does not address modeling or programming to the same extent.

The material on simulation modeling and programming, which is isolated to two chapters, uses Visual Basic for Applications (VBA) in Excel. This choice was driven by the reality that fewer and fewer graduate students come to me with programming experience. VBA/Excel is readily available, is easy to pick up quickly, and prepares students to learn Java, C++, or any other programming language later. This part of the book can be skipped for students who already know how to program simulations without compromising the remainder. Both Java and Matlab versions of Chap. 4, all of the software described in the book, and any data sets needed for exercises are available for download at the book website:

`users.iems.northwestern.edu/~nelsonb/IEMS435/`

The book should serve advanced undergraduates and graduate students. A prerequisite is a solid course in probability and statistics; statistics alone is not adequate. Although the book uses tractable stochastic process models (e.g., Markovian queues) as examples, the reader is not expected to have any background in these topics (in fact, students may find the stochastic processes course more intuitive and meaningful after having worked through this book). If students have had no experience with programming computer algorithms—for instance, in Matlab or some other programming language—then the instructor will have to supplement the book with more programming practice. Being an accomplished, or even good, programmer is not required, however.

Chapter 1 provides a concise summary of the book—except for programming—through a simple reliability problem. Chapter 2 then fills that gap by giving a quick start on simulation programming, using VBA. Chapter 3 introduces a number of tractable examples that illustrate the key issues that arise in analyzing stochastic systems via simulation; these examples recur throughout the book and so it is a must-read chapter. VBASim, a collection of VBA subs and class modules developed to support the book, is covered in Chap. 4 and is skippable if some other programming language or package will be used. Chapter 5 strengthens the connection between simulation and mathematical/numerical analysis of stochastic processes; it has the dual mission of setting up the design and analysis chapters that follow and preparing the student for more advanced courses on simulation methodology. Chapters 6–8 cover input modeling, output analysis, and experiment design, respectively, and are largely independent of the programming approach actually used to construct the simulation. The book concludes with Chap. 9, a guide to using simulation in research as opposed to using simulation to solve systems analysis problems.

What is missing? The book does not touch on the computing environment, and there are things you might want to do differently if you have, say, 500 CPUs available in a cloud computing cluster. I anticipate that by the time there is a need for a second edition, it will be easier to leverage such an environment for discrete-event, stochastic simulation, and I will add some general guidelines and recommendations. And while there is a lot of material on simulation optimization, the text is light on specific algorithms, reflecting the fact that there is no current agreement on baseline methods for all types of problems. That too will change. Finally, beyond discussing what it means, I did not do justice to the topic of validation of simulation models.

A number of students and colleagues contributed to the development of the programming approach used in this book. IEMS 435 initially used Law and Kelton (2000) as the text, and Dingxi Qiu spent a summer converting all of the C code in that book into VBA. Christine Nguyen helped in the development of VBASim, the simulation support library described in this book. Feng Yang worked on a research project with me where we used VBA for simulation analysis. Lu Yu assisted with the development of the solutions manual. The Java and Matlab versions of VBASim were translated by Luis de la Torre and Weitao Duan, respectively.

I have gotten a lot of feedback. Students in IEMS 435 suffered through incomplete versions of the text, spotting errors and typos with glee. Larry Leemis provided a thorough mark-up of an early draft, and Jason Merrick taught from it. It is good to have friends, and a number of mine read and marked up sections of the nearly complete book, including Christos Alexopoulos, Bahar Biller, John Carson, Xi Chen, Ira Gerhardt, Jeff Hong, Sheldon Jacobson, Seong-Hee Kim, Jack Kleijnen, Jeremy Staum, Laurel Travis, Feng Yang, Wei Xie, Jie Xu, and Enlu Zhou. Michael Fu guided my thinking about how to develop the section on gradient estimation, and Bruce Schmeiser did the same for error estimation. Seyed Iravani, Chuck Reilly, and Ward Whitt made sure I did not mangle the message of their papers in Chap. 9. In addition to those listed above, other people whose work influenced my think-

ing in this book include Sigrún Andradóttir, Russell Cheng, Dave Goldsman, Shane Henderson, David Kelton, Pierre L'Ecuyer, Lee Schruben, and Jim Wilson. Thank you all.

Last, but not least, work on this book was partially supported by National Science Foundation grant number CMMI-1068473. NSF does more with less than any federal government agency I know; let's keep it around.

Evanston, IL, USA

Barry L. Nelson

Contents

1	Why Do We Simulate?	1
1.1	An Example	1
1.2	Formalizing the Analysis	4
1.3	Issues and Extensions	5
	Exercises	5
2	Simulation Programming: Quick Start	7
2.1	A TTF Simulation Program	7
2.2	Some Important Simulation Concepts	11
2.2.1	Random-Variate Generation	11
2.2.2	Random-Number Generation	14
2.2.3	Replications	15
2.3	A VBA Primer	16
2.3.1	VBA: The Very Basics	16
2.3.2	Functions, Subs and Scope	21
2.3.3	Interacting with Excel	23
2.3.4	Class Modules	23
	Exercises	26
3	Examples	29
3.1	$M(t)/M/\infty$ Queue	29
3.2	$M/G/1$ Queue	31
3.3	The AR(1) Surrogate Model	34
3.4	A Stochastic Activity Network	35
3.5	Asian Option	37
	Exercises	39
4	Simulation Programming with VBASim	41
4.1	VBASim Overview	41
4.2	Simulating the $M(t)/M/\infty$ Queue	42
4.2.1	Issues and Extensions	48

4.3	Simulating the <i>M/G/1</i> Queue	49
4.3.1	Lindley Simulation of the <i>M/G/1</i> Queue	49
4.3.2	Event-Based Simulation of the <i>M/G/1</i> Queue	51
4.3.3	Issues and Extensions	57
4.4	Simulating the Stochastic Activity Network	58
4.4.1	Maximum Path Simulation of the SAN	58
4.4.2	Discrete-Event Simulation of the SAN	60
4.4.3	Issues and Extensions	61
4.5	Simulating the Asian Option	64
4.6	Case Study: Service Center Simulation	65
4.6.1	Issues and Extensions	73
	Exercises	75
5	Two Views of Simulation	79
5.1	A Framework for Simulation Modeling and Analysis	79
5.2	Simulation as a Stochastic Process	83
5.2.1	Simulated Asymptotic Behavior	83
5.2.2	Asymptotics Across Replications	88
5.2.3	Asymptotics Within Replications	90
5.2.4	Beyond Sample Means	92
	Appendix: Regenerative Processes and Steady State	94
	Exercises	95
6	Simulation Input	99
6.1	Input Modeling Overview	99
6.1.1	An Input Modeling Story	99
6.1.2	Characteristics of Input Processes	102
6.2	Univariate Input Models	104
6.2.1	Inference About Univariate Distributions	105
6.2.2	Estimation and Testing	109
6.2.3	Matching Properties of Univariate Distributions	113
6.2.4	Empirical Distributions	117
6.2.5	Direct Use of Input Data	119
6.2.6	Input Modeling Without Data	120
6.3	Multivariate Input Processes	122
6.3.1	Nonstationary Arrival Processes	123
6.3.2	Random Vectors	133
6.4	Generating Random Variates	134
6.4.1	Rejection Method	135
6.4.2	Particular Properties	137
6.5	Generating Pseudorandom Numbers	138
6.5.1	Multiple Recursive Generators	141
6.5.2	Combined Generators	141
6.5.3	Proper Use of Pseudorandom-Number Generators	143

- Appendix: Properties of the GLD 144
- Appendix: Implementation of the NORTA Method 145
- Exercises 149

- 7 Simulation Output** 157
 - 7.1 Performance Measures, Risk and Error 157
 - 7.1.1 Point Estimators and Measures of Error 158
 - 7.1.2 Measures of Risk and Error 161
 - 7.2 Input Uncertainty and Output Analysis 162
 - 7.2.1 Input Uncertainty: What Is It? 163
 - 7.2.2 Input Uncertainty: What to Do 164
- Appendix: Input Uncertainty for the $M/M/\infty$ Queue 167
- Exercises 168

- 8 Experiment Design and Analysis** 175
 - 8.1 Controlling Error via Replications 177
 - 8.2 Design and Analysis for Steady-State Simulation 181
 - 8.2.1 Fixed-Precision Problem 184
 - 8.2.2 Fixed-Budget Problem 187
 - 8.2.3 Batch Statistics 192
 - 8.2.4 Steady-State Simulation: A Postscript 193
 - 8.3 Simulation Optimization 194
 - 8.3.1 Convergence 197
 - 8.3.2 Correct Selection 198
 - 8.4 Experiment Design for Simulation Optimization 200
 - 8.4.1 Random-Number Assignment 201
 - 8.4.2 Design and Analysis for Correct Selection 206
 - 8.4.3 Adaptive Random Search 213
 - 8.4.4 Searching in Improving Directions 216
 - 8.4.5 Optimizing the Sample Average Problem 223
 - 8.4.6 Stochastic Constraints 225
 - 8.5 Variance Reduction Using Control Variates 229
 - 8.5.1 $M/G/1$ Queue Control Variate 233
 - 8.5.2 Stochastic Activity Network Control Variate 234
 - 8.5.3 Asian Option Control Variate 234
- Appendix: Properties of Control-Variate Estimators 235
- Exercises 236

- 9 Simulation for Research** 243
 - 9.1 Random Test Problems 244
 - 9.2 Robustness Studies 246
 - 9.3 Linking Experimental Factors 247
 - 9.4 Controlling Error in Research Simulations 249

- A VBASim** 253
 - A.1 Core Subs 253
 - A.2 Random-Number Generation 256
 - A.3 Random-Variate Generation 259
 - A.4 Class Modules 261
 - A.4.1 CTStat 261
 - A.4.2 DTStat 262
 - A.4.3 Entity 263
 - A.4.4 EventCalendar 263
 - A.4.5 EventNotice 263
 - A.4.6 FIFOQueue 264
 - A.4.7 Resource 265

- References** 267

- Index** 273

Chapter 1

Why Do We Simulate?

Stochastic simulation is a method for analyzing the performance of systems whose behavior depends on the interaction of random processes, processes that can be fully characterized by probability models. Stochastic simulation is a companion to mathematical and numerical analysis of stochastic models (e.g., Nelson 1995), and is often employed when the desired performance measures are mathematically intractable or there is no numerical approximation whose error can be bounded. Computers make stochastic simulation practical, but the method can be described independently of any computer implementation, which is what we do here.

1.1 An Example

We start with a simple, stylized example that nevertheless illustrates the key issues that we address in this book. We will use such examples throughout the text, and they have been carefully selected to exhibit the complexity of realistic problems, without being complicated to analyze or simulate. In both teaching and research, examples that illustrate complex behavior, but are not complicated to explain or analyze, are very important: they build intuition and provide a way to think through new ideas without obscuring things with a myriad of details. Realistic models may indeed be complicated, having many inputs and outputs and requiring thousands of lines of computer code to implement, so we also address modeling and programming simulations in Chaps. 2–4.

Example 1.1 (System time to failure TTF). The TTF system we consider here consists of two components that work as an active and a cold spare. The spare component becomes the active component when the (currently) active component fails, while the failed component immediately commences repair. The failed component becomes the spare when its repair is completed. Only one component at a time can be repaired, so the system fails if both components have failed, and it is operational as long as at least one of the components is functioning. The lifetime (time to

failure) of a component is equally likely to be 1, 2, 3, 4, 5 or 6 days, while repair takes exactly 2.5 days. Further, a repaired component is as good as new.

If we were interested in such a system, then we would probably want to know about its failure characteristics; for instance, its mean time to first system failure or its long-run system availability. Notice that while the failure characteristics of the components are completely specified (their lifetimes are equally likely to be from 1 to 6 days), the failure characteristics of the system as a whole—which depend on the interaction of failure and repair—are not immediately apparent. In fact, as simple as this system is it is difficult to derive these performance characteristics mathematically, while it is easy to simulate them (as we show below). This is why we simulate.

We will call those features of a system whose behavior is fully described by a probability model, such as the time to failure of a component, *inputs*. Derived quantities like the time of first system failure or system availability over some time horizon we call *outputs*.

The stochastic simulation method consists of generating realizations of the inputs, executing the system's logic to produce outputs and estimating system performance characteristics from the outputs.

Chapter 6 addresses representing and generating inputs; Chap. 7 covers analyzing the outputs, and Chap. 8 is about designing the experiment.

Let the *state* of the TTF system at any point in time be the number of functional components, 2, 1 or 0. The *events* that can cause the state of the system to change are the failure of a component and the completion of a repair. The current state, a list of pending future events—called an *event calendar*—and a *clock* are adequate to simulate *sample paths* of system behavior, provided we have a source of randomness to generate realizations of the inputs (time to failure of a component in this example). Here we will use a physical mechanism, a fair six-sided die, as our source of randomness.

Table 1.1 shows the sample path that would occur if the first four rolls of the die were 5, 3, 6, 1 and we stopped at the time of the first system failure (when the state reaches 0, meaning no functional components); each roll is inclosed in a box, so 6 represents a roll of the die yielding 6. The path was generated by executing the generic simulation algorithm shown in Fig. 1.1, noting that when a failure occurs the state decreases by one, while a repair increases it by one.

The approach illustrated in Table 1.1 often seems unnatural at first. For instance, why do we only keep track of the Next Failure when it is easy to figure out when each of the components will fail? And why not put the Next Repair on the calendar at time 0 since we can easily see it will be at time 7.5? The reason is the

Lazy simulator's rule: Schedule a future event only if not doing so could cause events to occur out of chronological order.

For instance, although we could figure out when both components will fail, we are not forced to schedule the failure of the spare until that moment when the active component has failed. Similarly, we do not have to schedule a repair to occur until

Table 1.1 Simulation of the TTF system until the first system failure

Clock	System state	Event calendar		Comment
		Next failure	Next repair	
0	2	$0 + \boxed{5} = 5$	∞	Initialize fully operational
5	1	$5 + \boxed{3} = 8$	$5 + 2.5 = 7.5$	One active, one under repair
7.5	2	8	∞	Next Failure remains on calendar
8	1	$8 + \boxed{6} = 14$	$8 + 2.5 = 10.5$	One active, one under repair
10.5	2	14	∞	Fully functional again
14	1	$14 + \boxed{1} = 15$	$14 + 2.5 = 16.5$	Active fails quickly
15	0	∞	16.5	System failure

Initialize: Set the initial values of the clock and system state, and put at least one future event on the event calendar.

Advance time: Update the clock to the time of the next pending event (and remove that event from the calendar).

Update: Update the state as appropriate for the current event and schedule any new future events to occur at the current clock time plus a time increment.

Terminate: If a termination condition has occurred, then stop; otherwise go to Advance time.

Fig. 1.1 Generic simulation algorithm

a component actually fails, even if we could work it out. On the other hand, at time 8 we *must* schedule both the Next Failure and Next Repair events because the evolution of the system from that point on depends on which of them occurs first. This rule turns out to be essential when programming complicated simulations that can have many types of events on the calendar simultaneously; in particular it avoids the need to cancel future events that become irrelevant because of another event that occurs first.

Here are some characteristics of the TTF example that are common to (nearly all) the stochastic simulations in this book:

- Simulated time (the simulation clock) jumps from event time to event time, rather than updating continuously; for this reason we call such simulations *discrete-event simulations*.
- The clock, the current state of the system, the list of future events and the event logic are all we need to advance the simulation to the next state change.
- The simulation ends when a particular system state is reached, at a fixed simulation time, or when a particular event occurs.

1.2 Formalizing the Analysis

Table 1.1 shows one sample path of the TTF system. Let the time of first system failure be denoted by Y , and the number of functional components at time t by $S(t)$. One sample path provides one observation of Y , but $S(t)$ is a function whose value evolves over the course of the simulation. If we are interested in the average number of functional components over some simulated time horizon T , denoted by $\bar{S}(T)$, then it is a *time average* because $S(t)$ has a value at all points in time. Thus,

$$\bar{S}(T) = \frac{1}{T} \int_0^T S(t) dt = \frac{1}{e_N - e_0} \sum_{i=1}^N S(e_{i-1}) \times (e_i - e_{i-1})$$

where $0 = e_0 \leq e_1 \leq \dots \leq e_N = T$ are the event times in the sample path. A characteristic of continuous-time outputs, like $S(t)$, in a discrete-event simulation is that they are piecewise-constant, since they can only change values at event times.

For the simulation in Table 1.1, the observed value of Y is 15, and the observed value of $\bar{S}(T)$ at $T = 15$ is

$$\begin{aligned} \frac{1}{15-0} [2(5-0) + 1(7.5-5) + 2(8-7.5) + 1(10.5-8) + 2(14-10.5) \\ + 1(15-14)] = \frac{24}{15}. \end{aligned}$$

Of course, this is only one possible pair of outputs $(Y, S(Y))$. *Replications* are statistically independent repetitions of the same model, and we will often make multiple replications to improve our estimates of system performance. An important distinction is between *within-replication* and *across-replication* output data. Both Y and $S(t)$ are *within-replication* outputs. The times of system failure Y_1, Y_2, \dots, Y_n , and the average number of functional components, $\bar{S}_1(T), \bar{S}_2(T), \dots, \bar{S}_n(T)$, from n different replications are *across-replication* outputs. Results across replications are naturally i.i.d. They are independent because we roll the die anew on each replication, and identically distributed because we apply the same initial conditions and model logic to those rolls of the die.

We run simulations to estimate system performance, often to compare alternative designs for a system. We can justify using simulation-based estimators when they satisfy some version of the *strong law of large numbers (SLLN)*, a topic we address more formally in Chap. 5. Two versions of the SLLN are relevant to stochastic simulation:

1. As the number of replications n increases, it may be that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n Y_i = \mu$$

(with probability 1), where μ can be interpreted as the the mean time to first system failure in the TTF example.

2. As the length of the replication T goes to infinity (which would make sense if we did not stop the simulation at the time of first system failure), it may be that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T S(t) dt = \theta$$

(with probability 1), where θ can be interpreted as the long-run average number of functional components in the TTF example.

The message is that as simulation effort increases (number of replications, length of replication, or perhaps both), the simulation estimators converge in a very strong sense to some useful system performance measures. This is comforting, but in practice we stop well short of infinity and thus do not fully converge. Therefore, an important topic is measuring the remaining error when we do stop; a strength of stochastic simulation is that we can do just that by using statistical inference.

1.3 Issues and Extensions

The TTF example illustrates the basics of stochastic simulation, but not all of the problems that can arise. Consider the following:

- Suppose that the repair times are not precisely 2.5 days, but instead are 1 day with probability 1/2, and 3 days with probability 1/2. A fair coin is all we need to generate repair times, but what if a failure and repair are scheduled to occur at the same time (which now can clearly happen)? Is it a system failure? Does the order in which we execute the event logic matter? Can we define a sensible tie-breaking rule?
- Suppose that there are three components rather than two. Is our definition of system state still adequate? What if more than one component can be repaired at a time? Do we need additional events?
- The long-run average number of functional components is not really “system availability.” Instead, we want the long-run fraction of time that at least one component is functional. How do we extract this performance measure from Table 1.1? How large should T be to get a good estimate of long-run availability?

Exercises

1. Have each member of the class simulate the TTF example independently until the time of first system failure. Estimate $E(Y)$, the expected value of the time to first system failure and also a standard error or confidence interval on this estimate.
2. Simulate the TTF example to the time of first system failure with three (one active, two spares) components rather than two.

3. Simulate the TTF example until time $T = 30$ and compute the average number of functional components and average system availability. Unless you are lucky, time 30 will not correspond to a failure or repair event time, so how can you stop the simulation right at time $T = 30$? *Hint:* The system is available if $S(t) > 0$. Average availability is the fraction of time the system is available from $t = 0$ to $t = T$.
4. For the TTF example, suppose that the repair times are not precisely 2.5 days, but instead are 1 day with probability $1/2$, and 3 days with probability $1/2$. Use a fair coin to generate repair times. Decide on a sensible tie-breaking rule and simulate until the time of the first system failure.
5. If in the TTF example the time to component failure and time to repair a component are exponentially distributed, then the TTF system is a continuous-time Markov chain for which the mean time to first failure and long-run system availability can be derived mathematically. If you know how, do this.
6. Write, in words, the logic for the two system events in the TTF example. Be sure to include updating the value of the state and scheduling any future events.
7. Add to Table 1.1 a column that updates the area $\sum_{i=1}^j S(e_{i-1}) \times (e_i - e_{i-1})$ when the j th event executes. Notice that by keeping this running total we can instantly calculate $\bar{S}(e_j)$ at any event time e_j .
8. Add to Table 1.1 a column that updates the area $\sum_{i=1}^j I\{S(e_{i-1}) = 2\} \times (e_i - e_{i-1})$ when the j th event executes, where I is the indicator function. Use this to compute the fraction of time that the system is fully functional (no component is in a failed state).
9. Consider a modified TTF system that works as follows. Simulate this system by hand (using dice) until the time of first system failure.
 - There are three components (one active, two spares, but still only one can be repaired at a time); repair time is 3.5 days.
 - In addition, each component is actually made up of two subcomponents, each subcomponent with a time to failure that is equally likely to be $1, 2, \dots, 6$ days.
 - A component fails when the *first* subcomponent fails. *In other words, to simulate the time to failure of a component, roll the die twice and take the smaller number.*
10. For your simulation in Exercise 9, add a column that updates the area $\sum_{i=1}^j S(e_{i-1}) \times (e_i - e_{i-1})$ when the j th event executes, where $S(t)$ is the number of functional components at time t . Use it to calculate the average number of functional components.

Chapter 2

Simulation Programming: Quick Start

As a first step toward more sophisticated simulation programming, this chapter presents a VBA simulation of the TTF example in Chap. 1. It also provides a gentle introduction to some important simulation concepts, and a brief overview of VBA, leaving the details to Chap. 4. Good reference texts on VBA as used for simulation and modeling are Elizandro and Taha (2008) and Albright (2007).

All readers should go through Sects. 2.1 and 2.2 of this chapter, even if they will not ultimately program their simulations in VBA. The focus here is on basic discrete-event simulation programming principles without using any special simulation support functions or objects. These principles form the core ideas for more sophisticated and scalable programming.

2.1 A TTF Simulation Program

Like all programming languages, VBA has variables. The key variables in the TTF simulation are S , representing the number of functioning components, `NextFailure`, which holds the time of the next pending component failure, `NextRepair`, the corresponding time of the next component repair, and `Clock`, the current time on the simulation clock.

The simulation will have a routine, called a `Sub` in VBA, for each event type. Three things happen in event routines:

1. The system state is updated;
2. Future events are scheduled; and
3. Statistics are computed or accumulated.

Figure 2.1 shows the event routines for the TTF simulation. We describe `Sub Failure` in some detail; the reader should work through the logic of `Sub Repair`.

When a failure occurs, the number of functional components always decreases by 1. If this leaves the system with only one functional component ($S = 1$), then

```

Private Sub Failure()
' Failure event
' Update state and schedule future events
  S = S - 1
  If S = 1 Then
    NextFailure = Clock + _
                  WorksheetFunction.Ceiling(6 * Rnd(), 1)
    NextRepair = Clock + 2.5
  End If

' Update area under the S(t) curve
  Area = Area + Slast * (Clock - Tlast)
  Tlast = Clock
  Slast = S
End Sub

Private Sub Repair()
' Repair event
' Update state and schedule future events
  S = S + 1
  If S = 1 Then
    NextRepair = Clock + 2.5
    NextFailure = Clock + _
                  WorksheetFunction.Ceiling(6 * Rnd(), 1)
  End If

' Update area under the S(t) curve
  Area = Area + Slast * (Clock - Tlast)
  Tlast = Clock
  Slast = S
End Sub

```

Fig. 2.1 Event routines for VBA simulation of the TTF problem

just prior to this event there were no failed components nor any component under repair. Therefore, we need to schedule the failure time of the spare component that has just become active, and schedule the repair completion for the active component that has just failed. Events are always scheduled at the current `Clock` time plus an increment into the future. Repairs always take 2.5 days, so `NextRepair = Clock + 2.5`. However, the time to failure of a component is random.

In VBA, the function `Rnd()` generates a random value that is approximately uniformly distributed between 0 and 1. Multiplying this number by 6 yields a value between 0 and 6; the `Ceiling` function rounds this up to an integer $\{1, 2, 3, 4, 5, 6\}$. Exactly why and how this works is a topic for later in the book.

The last part of each event routine is statistics updating. Recall that we showed in Chap. 1 that the time average number of functional components can be expressed as

$$\bar{S}(T) = \frac{1}{T} \int_0^T S(t) dt = \frac{1}{e_N - e_0} \sum_{i=1}^N S(e_{i-1}) \times (e_i - e_{i-1})$$

```

Private Function Timer() As String
    Const Infinity = 1000000

    ' Determine the next event and advance time
    If NextFailure < NextRepair Then
        Timer = "Failure"
        Clock = NextFailure
        NextFailure = Infinity
    Else
        Timer = "Repair"
        Clock = NextRepair
        NextRepair = Infinity
    End If
End Function

```

Fig. 2.2 Timer routine for VBA simulation of the TTF problem

where e_i are the event times. To avoid having to save all of the event times and values taken by the state variable, a standard trick is to keep a running sum of the area under the curve, updating it each time an event is executed. Specifically,

$$S(e_{i-1}) \times (e_i - e_{i-1}) = \text{Slast} * (\text{Clock} - \text{Tlast})$$

so that Clock is the time of the current event (e_i), Tlast is the time of the most recent prior event (e_{i-1}), and Slast is the value of $S(t)$ just before the change.

Figure 2.2 shows the VBA function used to advance the simulation from event to event. A function in VBA returns a single value, which in this case is the name of the next event to occur, either `Failure` or `Repair`. In addition, `Function Timer()` advances the simulation `Clock` to the time of the next event, and takes the triggering event off the event calendar by setting the time of its next occurrence to `Infinity`. For a simple simulation with only at most two pending events this approach is completely satisfactory; however, we will eventually need a method that scales up for more complicated simulations.

Figure 2.3 shows the global variable declarations and the program where execution of the simulation starts and ends. The `Dim` statement is used to declare variable names and their types; variables declared outside any `Sub` or `Function` are global variables whose values are visible and changeable by any routine in the program. It is clear that variables like `Clock`, `NextFailure` and `S` need to be global, since they are used in many places in the program.

`Sub TTF` is the main program that initializes and controls the execution of, and reports results from, the simulation. Typical initializations include setting up random-number generators and assigning initial values to key variables. Another critical initial step is scheduling at least one future event; since a discrete-event simulation jumps from event time to event time, without at least one pending event to start the simulation it will go nowhere.

```

Dim Clock As Double          ' simulation clock
Dim NextFailure As Double    ' time of next failure event
Dim NextRepair As Double     ' time of next repair event
Dim S As Double              ' system state
Dim Slast As Double          ' previous value of the system state
Dim Tlast As Double          ' time of previous state change
Dim Area As Double           ' area under S(t) curve

Private Sub TTF()
' Program to generate a sample path for the TTF example
  Dim NextEvent As String
  Const Infinity = 1000000
  Rnd (-1)
  Randomize (1234)

' Initialize the state and statistical variables
  S = 2
  Slast = 2
  Clock = 0
  Tlast = 0
  Area = 0

' Schedule the initial failure event
  NextFailure = WorksheetFunction.Ceiling(6 * Rnd(), 1)
  NextRepair = Infinity

' Advance time and execute events until the system fails
  Do Until S = 0
    NextEvent = Timer
    Select Case NextEvent
      Case "Failure"
        Call Failure
      Case "Repair"
        Call Repair
    End Select
  Loop

' Display output
  MsgBox ("System failure at time " _
    & Clock & " with average # functional components " _
    & Area / Clock)
  End
End Sub

```

Fig. 2.3 Declarations and main program for VBA simulation of the TTF problem

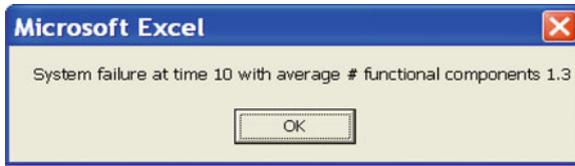


Fig. 2.4 Output from the TTF simulation

The main program also contains a loop that repeatedly advances time to the next pending event and executes that event until some ending condition occurs. The TTF simulation continues until $S = 0$, meaning the system fails (no functioning components) for the first time. Simulations can also end when a particular event occurs, as we will see in later examples.

The final section of the main program reports two output results: the time of system failure and the average number of functional components from time 0 until the system first fails. Notice that since the simulation ends as soon as $S = 0$, the simulation `Clock` is also the time of system failure. And since we have been accumulating the area under the $S(t)$ curve in the variable `Area`, `Area/Clock` is the average value of $S(t)$ up to that time. The output from the simulation is displayed in Fig. 2.4.

2.2 Some Important Simulation Concepts

This chapter is a “quick start” to simulation programming; to complete it, three additional topics need more explanation: Random-variate generation, random-number generation, and replications

2.2.1 Random-Variate Generation

A stochastic simulation is driven by one or more fully characterized probability models from which we can generate realizations. We use the term “realization” to mean specific, usually numerical, outcomes, and we call the process of creating these realizations *random-variate generation*. In simulation we do not manipulate symbols, as in a mathematical analysis, we generate realizations—actual numerical values—that are consistent with (in a way we define precisely below) a probability model. This makes our ultimate analysis statistical, not mathematical, since we *estimate* performance measures of interest from simulated data.

The most basic driving probability models we need are i.i.d. random variables X_1, X_2, \dots with known distribution F_X . In the TTF simulation we need failure times that are i.i.d. with a discrete uniform distribution on $\{1, 2, \dots, 6\}$; in Chap. 3 we introduce examples requiring other probability models, including the exponential distribution.

Suppose we have an algorithm called `algorithm`, which takes a single numerical input u and produces a single numerical output x ; that is

$$x = \text{algorithm}(u).$$

Here `algorithm` is typically some computer code, but the important thing is that it is a deterministic input–output transformation. For instance, it might be that `algorithm(0.3)` returns the value $x = 2$.

Now suppose that the input u is not specified, but instead is a random variable, say U , that has a known probability distribution F_U . Then `algorithm` defines a random variable

$$X = \text{algorithm}(U)$$

and it makes perfect sense to talk about $\Pr\{X \leq 20\}$, or the $E(X)$, or even its distribution F_X . The important concept here is that `algorithm`, which we think of as a recipe that takes a numerical input and yields a numerical output, can also be thought of as a way to transform a random variable U with some distribution into a random variable X with a different distribution. A random variable is a mathematical object about which we can make probability statements. In a very real sense, all of stochastic simulation is built on this insight.

A very important special case is when F_U is the uniform distribution on the interval $[0, 1]$, meaning

$$F_U(u) = \Pr\{U \leq u\} = \begin{cases} 0, & u < 0 \\ u, & 0 \leq u < 1 \\ 1, & 1 \leq u. \end{cases} \quad (2.1)$$

This distribution is often denoted $U(0, 1)$.

As a particular case of `algorithm` suppose `algorithm(u) = -VBA.Log(1 - u)`, where `VBA.Log` is the VBA function that returns the natural logarithm of its argument. Assuming the `VBA.Log` function works perfectly, what is the distribution of the value returned when U is $U(0, 1)$?

$$\begin{aligned} \Pr\{\text{algorithm}(U) \leq x\} &= \Pr\{-\text{VBA.Log}(1 - U) \leq x\} \\ &= \Pr\{1 - U \geq e^{-x}\} \\ &= \Pr\{U \leq 1 - e^{-x}\} \\ &= 1 - e^{-x}, \quad x \geq 0. \end{aligned} \quad (2.2)$$

The last step follows because $\Pr\{U \leq u\} = u$ for $0 \leq u \leq 1$.

Expression (2.2) is the cumulative distribution function (cdf) of the exponential distribution with mean 1, implying we have created an algorithmic description of

this random variable. Thus, if we had a way to produce realizations of U , we could generate realizations of the exponential with mean 1; this is random-variate generation.

Now, rather than wondering what F_X we get when we apply some algorithm to U , we want to *construct* algorithm so that X has a distribution F_X that we specify. Being able to do this is absolutely essential to stochastic simulation.

Consider first the case where the cdf F_X is a continuous, increasing function for all x such that $0 < F_X(x) < 1$; this is true for the exponential, normal, lognormal, Weibull, uniform, and many other well-known distributions. Under these conditions the equation $u = F_X(x)$ has a unique solution which we denote by $x = F_X^{-1}(u)$, the *inverse cdf*. Then for $U \sim U(0, 1)$, define the random variable X by $X = F_X^{-1}(U)$. We can immediately show that

$$\begin{aligned} \Pr\{X \leq x\} &= \Pr\{F_X^{-1}(U) \leq x\} \\ &= \Pr\{U \leq F_X(x)\} \\ &= F_X(x). \end{aligned}$$

This is how we obtained algorithm for the exponential distribution with mean 1 for which $F_X^{-1}(u) = -\ln(1-u)$. Even when F_X^{-1} is not available in closed form, as it is for the exponential, the cdf can be inverted numerically using a root-finding algorithm. Thus, the inverse cdf method is a general-purpose method for distributions with continuous, increasing cdfs.

If we define the inverse a bit more carefully, then the inverse cdf method works for any random variable:

$$X = F_X^{-1}(U) = \min\{x : F_X(x) \geq u\}. \quad (2.3)$$

When F_X is continuous and increasing, the two definitions coincide. But (2.3) also works for cdfs that contain jumps.

In particular, suppose that X can take values $x_1 < x_2 < x_3 < \dots$ with corresponding probabilities p_1, p_2, p_3, \dots , where $\sum_i p_i = 1$. The time to component failure in the TTF system is of this type, where $x_i = i$ and $p_i = 1/6$ for $i = 1, 2, \dots, 6$. Then applying the Definition (2.3)

$$\begin{aligned} \Pr\{X = x_j\} &= \Pr\{F_X^{-1}(U) = x_j\} \\ &= \Pr\{F_X(x_j) \geq U > F_X(x_{j-1})\} \\ &= \Pr\{F_X(x_{j-1}) < U \leq F_X(x_j)\} \\ &= \Pr\{U \leq F_X(x_j)\} - \Pr\{U \leq F_X(x_{j-1})\} \\ &= F_X(x_j) - F_X(x_{j-1}) \end{aligned} \quad (2.4)$$

$$\begin{aligned} &= \sum_{i=1}^j p_i - \sum_{i=1}^{j-1} p_i \\ &= p_j \end{aligned} \quad (2.5)$$

as desired. Expression (2.4) provides the basic algorithm for discrete distributions: Given U , find the smallest j such that $\sum_{i=1}^j p_i \geq U$ and return x_j .

As an illustration, consider the discrete uniform distribution on $\{1, 2, \dots, k\}$, of which the time to component failure in the TTF example is a special case with $k = 6$. Thus, $x_i = i$ and $p_i = 1/k$ for $i = 1, 2, \dots, k$. Expression (2.4) implies we need to find j such that

$$\frac{j-1}{k} < U \leq \frac{j}{k}. \quad (2.6)$$

In the exercises you are asked to show that this leads to the formula $X = \lceil Uk \rceil$, where $\lceil \cdot \rceil$ is the ceiling function that returns the smallest integer that is greater than or equal to its argument.

We say that a random-variate generation algorithm is exact if, under the assumption of $U(0, 1)$ variates as input, and infinite precision for the numerical variables and functions that make up algorithm, it can be proven to imply the desired distribution using an argument such as (2.2).

Thus, our algorithms for the exponential distribution and discrete uniform distribution are exact.

In practice there are compromises. The first is that functions such as `VBA.Log` do not work perfectly, if for no other reason than the computer representation of numbers is finite. This is an issue in all numerical programming. The second compromise is more central to simulation and has to do with generating realizations of U .

2.2.2 Random-Number Generation

Producing realizations of U is called *random-number generation*. The definition of exact random-variate generation for X requires that $U \sim U(0, 1)$; therefore, an exact method to generate X_1, X_2, \dots that are i.i.d. F_X requires U_1, U_2, \dots that are i.i.d. $U(0, 1)$. Because of the finite representation of numbers in a computer we know that this is not strictly possible. Further, even if it were possible, it is a bad idea, since the simulation program would produce different output each time it is run, even if nothing was changed.

Instead, stochastic simulations use *pseudorandom numbers*. The easiest way to visualize this is to think of a large, ordered list of numbers between 0 and 1:

$$U_1, U_2, U_3, \dots, U_i, U_{i+1}, \dots, U_{P-1}, U_P, U_1, U_2, \dots$$

The list needs to be quite long (about two billion in older simulation programs, much, much longer in newer ones). There may be repeats in this list, but when the end of the list is reached it starts over from the beginning; for this reason P is called the *period* of the pseudorandom numbers. A good list of pseudorandom numbers appears to be statistically indistinguishable from random numbers, provided the portion you use in your simulation is significantly less than P .

Obviously, pseudorandom numbers are not random in any sense: they are a pre-determined, finite list that will eventually repeat. The VBA code

```
Rnd (-1)
Randomize (1234)
```

sets the starting point in the list, and the VBA function `Rnd()` draws them out in order from that starting point. This is why every time the TTF program is run the same output results. If 1234 is replaced by a different integer, then the pseudorandom numbers would be drawn starting from a different place in the list and thus the result would be different (you can try this).

Actually storing a list of pseudorandom numbers would be cumbersome at best, and impossible if P is large enough. Therefore, the pseudorandom numbers in simulation software are actually produced by a recursive algorithm; we describe the construction of such pseudorandom-number generators in Chap. 6, but the list representation is conceptually correct. Using a pseudorandom-number generator, rather than random numbers, is an important practical compromise whose full implications are revealed throughout the book.

2.2.3 Replications

As discussed in Chap. 1, one justification for using simulation-based estimators is that as the number of replications increases the sample average across replications converges (with probability 1) to the long-run average. In this section we introduce the key ideas behind programming replications.

To be concrete, consider both the random variable Y , the time to system failure, and $\bar{S}(Y)$, the average number of functional components up to time Y , in the TTF example. Suppose we want to simulate 100 i.i.d. replications of each of these, and report their overall averages. Specifically, we want

$$\bar{Y} = \frac{1}{100} \sum_{i=1}^{100} Y_i$$

$$\bar{S} = \frac{1}{100} \sum_{i=1}^{100} \bar{S}_i(Y_i)$$

where the subscript i represents the results from the i th replication.

To constitute “replications” Y_1, Y_2, \dots, Y_{100} should be i.i.d., and the same should be true for $\bar{S}_1(Y_1), \bar{S}_2(Y_2), \dots, \bar{S}_{100}(Y_{100})$. Identically distributed means that the simulation logic, including the initial conditions, are the same on each replication. Independent means that the component failure times on any replication are statistically independent of the failure times on any other replication. The failure times depend on the random numbers, and as discussed above, the sequence of random numbers produced by a pseudorandom-number generator approximate i.i.d. behavior, provided we just let the generator keep running and do not start over. Thus, the

replications will be (approximately) i.i.d. provided we *do not* reset the pseudorandom-number generator between replications, but we *do* reset everything else.

Figure 2.5 shows a modification of the TTF simulation's main program to implement replications. The pseudorandom-number generator, and the across-replication statistical variables, are initialized once outside the replication loop `For Rep = 1 To 100`. The state variable `S`, the within-replication statistical variables, and the initial failure and repair events are initialized inside the replication loop. In this way we achieve (approximately) i.i.d. results. A key to simulation experiment design is choosing the number of replications large enough so that where we start the pseudorandom-number generator does not matter.

Notice how at the end of each replication we update sums of the average number of functional components and the time of system failure, so that when all of the replications are complete, we can easily compute the across-replication averages. This is an alternative to saving the results from each replication and computing the overall average later. Each approach has its advantages and disadvantages. The result of running this simulation is shown in Fig. 2.6.

2.3 A VBA Primer

In this book we illustrate simulation programming using the Visual Basic for Applications (VBA) language in Excel, along with some custom support routines called VBASim. VBASim and a number of simulation programming examples are presented in Chap. 4. VBA for Excel was chosen because it is widely available (if you have Excel, you have VBA), relatively easy to learn, and Excel itself provides a large collection of useful functions as well as worksheets on which to write simulation output data. The Preface to the book describes how to obtain Java or Matlab versions of the TTF example and all of the simulations in Chap. 4 if you already know Java or Matlab.

This section presents a brief introduction to VBA. Here and in Chap. 4 we have avoided relying on or describing any features of Excel that are likely to change such as specific interface, menu or formatting options. If you are learning VBA and Excel from the beginning and entirely on your own, then you will likely need a comprehensive reference such as Walkenbach (2010) in addition to this section.

2.3.1 VBA: *The Very Basics*

VBA code is organized into modules or class modules. Think of these as folders. In the simplest cases, such as the TTF example, there is a single Module that contains all of the VBA code for the simulation. We will assume that for now.

```

Private Sub TTFRep()
' Program to generate a sample path for the TTF example
  Dim NextEvent As String
  Const Infinity = 1000000
  Rnd (-1)
  Randomize (1234)

' Define and initialize replication variables
  Dim Rep As Integer
  Dim SumS As Double, SumY As Double
  SumS = 0
  SumY = 0

  For Rep = 1 To 100

' Initialize the state and statistical variables
    S = 2
    Slast = 2
    Clock = 0
    Tlast = 0
    Area = 0

' Schedule the initial failure event
    NextFailure = WorksheetFunction.Ceiling(6 * Rnd(), 1)
    NextRepair = Infinity

' Advance time and execute events until the system fails
    Do Until S = 0
      NextEvent = Timer
      Select Case NextEvent
        Case "Failure"
          Call Failure
        Case "Repair"
          Call Repair
      End Select
    Loop

' Accumulate replication statistics
    SumS = SumS + Area / Clock
    SumY = SumY + Clock
  Next Rep

' Display output
  MsgBox ("Average failure at time " _
    & SumY / 100 & " with average # functional components " _
    & SumS / 100)

  End
End Sub

```

Fig. 2.5 Main program for VBA simulation of the TTF problem with replications

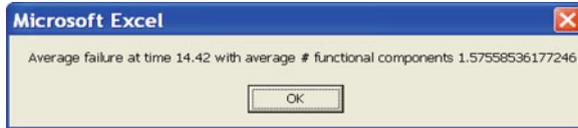


Fig. 2.6 Output from the TTF simulation across 100 replications

A module begins with variable declarations, followed by one or more subs and functions. The subs and functions contain the executable VBA code that describes the simulation.

VBA has a number of variable types; the ones we will need are

Double: double-precision real values
 Integer: integer values in the range $-32,768$ – $32,768$
 Long: integers in the range of roughly ± 2 billion
 String: character variables
 Boolean: variables with only True or False values

Variable types are declared using one of the keywords `Dim`, `Private` or `Public`. We will defer describing `Public` until later; `Dim` and `Private` mean the same thing and we will typically use `Dim` in this book. The declarations at the beginning of the TTF module are

```
Dim Clock As Double           ' simulation clock
Dim NextFailure As Double     ' time of next failure event
Dim NextRepair As Double      ' time of next repair event
Dim S As Double               ' system state
Dim Slast As Double           ' previous value of the system state
Dim Tlast As Double           ' time of previous state change
Dim Area As Double            ' area under S(t) curve
```

Variable names may be lengthy and can contain letters and numbers, but not spaces. VBA recognizes any capitalization you use in the declaration and forces you to use it throughout the program. For instance, if you tried to use the variable `area` in the TTF program then VBA would automatically change it to `Area`. More than one variable may be declared in a single statement, as in

```
Dim Code As Integer, Product As String, InStock as Boolean
```

It is worth noting here that when VBA sees a single quote `'` it treats all text from that point to the end of the current line as a nonexecutable comment.

Variables declared `Dim` or `Private` at the beginning of the module are visible to all subs or functions within that module. Variables declared `Dim` or `Private` within a sub or function are visible only within that sub or function, and only retain

their values while that sub or function is being executed. Consider the following portion of the replication version of the TTF program:

```
Private Sub TTFRep()
' Program to generate a sample path for the TTF example
  Dim NextEvent As String
  Const Infinity = 1000000
  Rnd (-1)
  Randomize (1234)
' Define and initialize replication variables
  Dim Rep As Integer
  Dim SumS As Double, SumY As Double
```

The variables `NextEvent`, `Rep`, `SumS` and `SumY` are declared within `Sub TTFRep` so their values are not visible to other subs or functions within the module; in fact, these variable names could be used in other subs or functions without conflict.

VBA also allows multidimensional arrays, declared as shown in this illustration:

```
Dim X(1 to 10) as Double, Y(0 to 5, 0 to 20) as Integer
```

Notice that array indices need not start at 1.

We will frequently use three control structures: `If-Then-Else-EndIf`, `Select Case` and various kinds of loops. The generality of `If-Then-Else-EndIf` is illustrated below:

```
If Index = 0 Then
  X = X + 1
  Y = VBA.Sqrt(X)
Else If Index = 1 Then
  Y = VBA.Sqrt(X)
Else If Index = 2 Then
  Y = X
Else
  X = 0
End If
```

Control structures in VBA always have an explicit ending statement (`If` is ended by `End If` here).

The `Select Case` structure is an alternative to `If` when there are a number of possible outcomes based on the value of a single variable. In the following example the variable `Index` is compared to a single value (0), a range (1 To 10), a condition (`Is < 0`), another variable (`MyIndex`) and none of the above (`Else`). The statements that follow the first match are executed.

```
Select Case Index
  Case 0
    [VBA statements...]
  Case 1 to 10
    [VBA statements...]
  Case Is < 0
    [VBA statements...]
  Case MyIndex
    [VBA statements...]
```

```

    Case Else
        [VBA statements...]
    End Select

```

In the TTF simulation the `Select Case` structure is used to execute the appropriate event sub depending on the value of the variable `NextEvent`.

Loops are used to execute a group of statements repeatedly. Three common forms are

```

For counter = start To end [Step increment]
    [VBA statements...]
Next counter

```

```

Do
    [VBA statements...]
Loop [While|Until] condition

```

```

Do [While|Until] condition
    [VBA statements...]
Loop

```

The `For-Next` loop increments the `counter` variable in steps of size `increment` (the default is 1) until the `end` value is met or exceeded. The two forms of the `Do-Loop` structure differ only in whether the ending condition is tested before or after the first execution of the loop. The following piece of `Sub TTFRep()` illustrates both types of loops. Notice also the use of the `Select Case` structure to execute the next event.

```

    For Rep = 1 To 100
' Initialize the state and statistical variables
        S = 2
        Slast = 2
        Clock = 0
        Tlast = 0
        Area = 0
' Schedule the initial failure event
        NextFailure = WorksheetFunction.Ceiling(6 * Rnd(), 1)
        NextRepair = Infinity
' Advance time and execute events until the system fails
        Do Until S = 0
            NextEvent = Timer
            Select Case NextEvent
                Case "Failure"
                    Call Failure
                Case "Repair"
                    Call Repair
            End Select
        Loop
' Accumulate replication statistics
        SumS = SumS + Area / Clock
        SumY = SumY + Clock
    Next Rep

```

VBA provides a graceful way to exit any control structure before it has completed via the `Exit` statement. For instance, if an `Exit For` or `Exit Do` statement is encountered in a `For` or `Do` loop, respectively, it causes the loop to terminate immediately without checking the ending condition.

2.3.2 *Functions, Subs and Scope*

Execution of a VBA program takes place in functions and subs. Functions allow arguments to be passed as inputs, and return a single value as output via the function name. Consider Function `Timer`:

```
Private Function Timer() As String
    Const Infinity = 1000000
    ' Determine the next event and advance time
    If NextFailure < NextRepair Then
        Timer = "Failure"
        Clock = NextFailure
        NextFailure = Infinity
    Else
        Timer = "Repair"
        Clock = NextRepair
        NextRepair = Infinity
    End If
End Function
```

`Timer` has no arguments since the variables `Clock`, `NextFailure` and `NextRepair` are visible to the entire module; it returns a single value of type `String`. The type of the returned value must be declared in the definition of the Function. The statement

```
NextEvent = Timer
```

calls the function, returning the name of the next event to be executed. `Timer` also illustrates another data type, called a constant (`Const`), a declaration that assigns a fixed unchangeable value to a variable name.

A sub, on the other hand, does not (directly) return a value and therefore does not have a declared type; instead, it is a subprogram that performs some calculations. For instance, the statement `Call Failure` causes the following sub to be executed, and then returns control to the next statement following the call.

```
Private Sub Failure()
    ' Failure event
    ' Update state and schedule future events
    S = S - 1
    If S = 1 Then
        NextFailure = Clock + WorksheetFunction.Ceiling
            (6*Rnd(), 1)
        NextRepair = Clock + 2.5
    End If
    ' Update area under the S(t) curve
    Area = Area + Slast * (Clock - Tlast)
```

```

    Tlast = Clock
    Slast = S
End Sub

```

A graceful way to leave a sub or function before reaching the `End Sub` or `End Function` statement is to use, respectively, the `Exit Sub` or `Exit Function` statement.

In many contexts it is important to be able to pass arguments (inputs) to functions and subs; this allows them to be more flexible. For instance, in Chap.4 we will introduce `Sub Report` that is useful for writing simulation results to an Excel worksheet. Here is the entire sub:

```

Public Sub Report(Output As Variant, WhichSheet As String, _
    Row As Integer, Column As Integer)
    Worksheets(WhichSheet).Cells(Row, Column) = Output
End Sub

```

The sub has four arguments: the output value to be written, the name of the Excel worksheet to which to write it, and the row and column numbers of the cell to which it should be written. Notice that the type of each argument is declared in the definition of the Sub. Type `Variant` is used for a variable whose type is not known in advance, which is relevant here since the output could be an integer, real or even string. Also notice that the underscore character `_` is used to let VBA know that a statement continues on the next line. `Sub Report` might be called as follows:

```

Call Report(Clock, "Sheet1", 2, 1)

```

This would write the current value of the variable `Clock` to the second row and first column (cell A2) of the Excel worksheet `Sheet1`.

An important fact about arguments to functions and subs is that if you change the value of any argument within the function or sub then the changed values are also changed in the calling sub or function. In VBA terminology, the arguments are passed “by reference.” In this way a sub can indirectly return a value. To avoid this, arguments can be passed “by value.” For instance, we could rewrite the first line of `Sub Report` to be

```

Public Sub Report(ByVal Output As Variant, _
    ByVal WhichSheet As String, _
    ByVal Row As Integer, ByVal Column As Integer)

```

You probably noticed the declarations `Public` and `Private` in the definitions of the subs and functions. A sub or function that is declared `Private` can only be called within the module that contains its code. All of the code for the TTF simulation is in the same module, so this is fine. A sub or function that is declared `Public` can be called from any module in the same Excel workbook. All VBASim subs and functions are `Public` so that you do not need to include them in the module containing your simulation; instead they can reside in the VBASim module in the same workbook.

The visibility of a sub or function is referred to as its “scope.” As we have seen, variables also have scope. To allow a variable to be visible to other modules in

the same workbook it must be declared `Public` at the top of its module. Many VBASim variables are `Public`.

2.3.3 Interacting with Excel

We will interact with Excel in two ways: reading values from and writing values to an Excel worksheet, and using Excel functions within VBA programs.

There are a number of ways to read from and write to Excel worksheets, but the most basic is the `.Cells` method, which is used in `Sub Report` of the previous section. To use this method we need to know the name of the worksheet, and the row and column to write to or read from. For instance, to write the value of `Clock` to row 2 and column 20 of `Sheet1`, we use the statement

```
Worksheets("Sheet1").Cells(2,20) = Clock
```

Similarly, to read the number of call center agents from row *I* and column *J* of the worksheet `Fax` we use

```
NumAgents = Worksheets("Fax").Cells(I, J)
```

VBA has a small number of built-in intrinsic functions; these include square root (`VBA.Sqrt`), natural logarithm (`VBA.Log`) and absolute value (`VBA.Abs`). Excel, on the other hand, has a vast number of intrinsic functions. To use them in VBA, the prefix `WorksheetFunction.` is added to the name of the Excel function. For instance, to use the Excel function `Max` that returns the maximum of a list of arguments *W, X, Y, Z*, the statement would be

```
M = WorksheetFunction.Max(W, X, Y, Z)
```

2.3.4 Class Modules

VBA supports a basic level of object-oriented programming; if you are already familiar with object-oriented programming in languages such as Java and C++, then the approach in VBA will seem familiar. This section provides a quick introduction to this very important topic.

Discrete-event simulations often contain multiple instances of similar objects:

- *Entities*: things that come and go, such as customers, messages, jobs, events
- *Queues*: ordered lists of entities
- *Resources*: scarce quantities that entities need, including servers, computers or machines
- *Statistics*: information to be recorded on all of the above

User-defined objects in VBA are called class modules, and a class module is the template for an object. Class modules contain

- *Properties*: variables (often called “attributes” in simulation terminology) that can retain values
- *Methods*: instructions (subs and functions) about how to do things

The key benefit of objects is that we can create multiple instances from the same template, each uniquely identified and with its own properties and methods.

To illustrate the point, recall the continuous-time statistic average number of functional components we maintained in the TTF simulation. To do so required defining and updating three variables: `Slast`, the previous value of the system state; `Tlast`, the time of the previous state change; and `Area`, the area under the $S(t)$ curve. If we had additional continuous-time statistics (e.g., average system availability) then we would need a corresponding set of variables and updates. In a realistic simulation we might have many such statistics.

Figure 2.7 shows the VBASim class module `CTStat` that is used to record and compute continuous-time statistics. In a VBA program where this class module is present, the two statements

```
Dim Functional As New CTStat
Dim Available As New CTStat
```

would cause two unique `CTStat` objects to be created based on this template. Then, for instance, the method

```
Functional.Record(S)
```

could be used to update the area under the $S(t)$ curve when the value of S changes, while the method

```
Abar = Available.Mean
```

would compute the average availability up to the current time. Each object maintains its own unique `Tlast`, `Xlast` and `Area` properties without conflict. We now describe the `CTStat` class module itself in a little more detail.

The declarations in the beginning define the properties that the object has. These should be `Private` if they are only used within the object and we want each instance of the object to have its own unique copy. The methods (subs and functions) are accessed by adding the unique object name as a prefix to the name of the sub or function, as in `Functional.Record(S)` or `Available.Mean`. A class module may optionally have a special sub called `Private Sub Class_Initialize()` that is executed whenever a new instance of the object is created.

VBA provides a special type of data structure, called a collection, that can store objects in much the same way that an array stores numerical or character data. To illustrate, suppose that we had a large number of `CTStat` objects. Then we might define a collection by

```
Dim TheCTStats As New Collection
```

All Collections have an `Add` method for inserting an object; for instance

```
TheCTStats.Add Available, Before:=2
```

```

' Generic continuous-time statistics object
' Note that CTStat should be called AFTER the value
' of the variable changes

Private Area As Double
Private Tlast As Double
Private Xlast As Double
Private TClear As Double

Private Sub Class_Initialize()
' Executes when CTStat object is created to initialize variables

    Area = 0
    Tlast = 0
    TClear = 0
    Xlast = 0
End Sub

Public Sub Record(X As Double)
' Update the CTStat from last time change and keep track
' of previous value

    Area = Area + Xlast * (Clock - Tlast)
    Tlast = Clock
    Xlast = X
End Sub

Function Mean() As Double
' Return the sample mean up through current time
' but do not update

    Mean = 0
    If (Clock - TClear) > 0 Then
        Mean = (Area + Xlast * (Clock - Tlast)) / (Clock -
            TClear)
    End If
End Function

Public Sub Clear()
' Clear statistics

    Area = 0
    Tlast = Clock
    TClear = Clock
End Sub

```

Fig. 2.7 CTStat class module from VBASim

inserts the `CTStat Available` before the object currently in position 2; there is also an `After` option. Similarly, the `TheCTStats.Remove(j)` method removes the object currently in position `j`; the `TheCTStats.Item(i)` method points to the object currently in position `i`; and the `TheCTStats.Count` method returns the number of objects currently in the collection.

Exercises

1. Modify the TTF simulation so that it can have any number of components, not just 2. Assume that components can still be repaired only one at a time.
2. Simulate the TTF example until time $T = 1,000$ and report the average number of functional components. To get the simulation to stop at time 1,000, create a new event called `EndSimulation`, with a variable `NextEndSimulation` as part of the event calendar. Now `Timer()` will chose the next event from among `NextFailure`, `NextRepair` and `NextEndSimulation`. Why is this a more effective approach than changing the loop condition to `Do Until Clock >= 1000`?
3. For the TTF simulation, suppose that the repair times are NOT precisely 2.5 days, but instead are 1 day with probability $1/2$, and 3 days with probability $1/2$. Decide on a sensible tie-breaking rule, modify the simulation and simulate until the time of the first system failure.
4. The TTF system is available provided there is at least one functional component. Define a variable $A(t)$ that takes the value 1 when the system is available, and is 0 otherwise. Then

$$\bar{A}(T) = \frac{1}{T} \int_0^T A(t) dt$$

is the average system availability from time 0 to T . Modify the simulation created for Exercise 2 to also report average system availability.

5. For each of the exercises above, add the capability to do replications and write the outputs from each replication to an Excel worksheet, one row per replication. Using 100 replications, estimate the expected value of each output and a 95% confidence interval on it.
6. Show that (2.6) implies that $X = \lceil Uk \rceil$ is the inverse cdf for the discrete uniform distribution on $\{1, 2, \dots, k\}$.
7. An alternative to using $X = \lceil Uk \rceil$ to generate variates from a discrete uniform distribution on $\{1, 2, \dots, k\}$ is $X = \lfloor Uk + 1 \rfloor$. Both methods have an flaw if it is possible for $U = 0$ or $U = 1$. What is it? For this reason (and others) pseudorandom-number generators are designed so that they do not generate exactly 0 or 1.
8. Derive a method for generating a discrete uniform distribution on $\{a, a + 1, a + 2, \dots, b\}$ where a and b are integers with $a < b$.

9. Modify the VBA code for the TTF simulation with replications to match the description from Exercise 9 of Chap. 1. Write the results from each replication (average number of functional components and the time to failure) to one row of an Excel worksheet. Using 100 replications, estimate the expected value of each performance measure and a 95% confidence interval on it.
10. A very simple inventory system works as follows. At time 0 it has 50 items in stock. Each day there is a demand for items which is equally likely to be 1, 2, ..., 9 items. The system satisfies as much of the demand as possible, but if it does not have enough items in stock then the excess demand is lost. When the stock reaches or goes below `ReorderPoint`, an order is place for 50 additional items; it takes 1, 2 or 3 days for the order to arrive, equally likely. There can only be one pending order at a time, meaning that if an order has been placed and has not yet arrived then another order cannot be placed. Orders arrive before demands each day, and the planning horizon is 365 days. Build a simulation model of this system, and use it to find the smallest value of `ReorderPoint` that guarantees that there are no items in stock less than 5% of the time. Also estimate the mean stock level for your policy. Put a 95% confidence interval around your estimate. Start with 100 replications, but increase this as necessary to be sure of your recommendation.

Chapter 3

Examples

None of the examples in this chapter need to be simulated; they can all be analyzed by mathematical/numerical analysis. But thinking about how to simulate them will help us understand what works, and why, for systems we do need to simulate. These examples also provide a tested for evaluating new ideas in simulation design and analysis. You may have encountered these models in other classes, but it is neither necessary to have seen them before, nor to know how the results are derived, to use them throughout the book.

3.1 $M(t)/M/\infty$ Queue

This example is based on Nelson (1995, Chap. 8).

Example 3.1 (The parking lot). Prior to building a large shopping mall, the mall designers must decide how large a parking garage is needed. The arrival rate of cars will undoubtedly vary over time of day, and even day of year. Some patrons will visit the mall for a brief time, while others may stay all day. Once the garage is built, the mall may open and close floors depending on the load, but the first question is, what should the maximum capacity of the garage be?

Since the mall developers would like virtually everyone to be able to have a parking space, a standard modeling trick is to pretend that the parking garage is infinitely large and then evaluate the probability distribution of the number of spaces actually used. If the capacity of the garage is set to a level that would rarely be exceeded in an infinitely large garage, then that should be adequate in practice.

The $M(t)/M/\infty$ queue is a service system in which customers arrive according to a nonstationary Poisson arrival process with time-varying arrival rate $\lambda(t)$ customers/time,¹ the service time is exponentially distributed with mean τ , and there

¹ A stationary Poisson arrival process has independent times between arrivals that are exponentially distributed with mean $1/\lambda$ or equivalently arrival rate λ . The nonstationary Poisson process is a generalization that allows the arrival rate to change over time; it is covered in detail in Chap. 6.

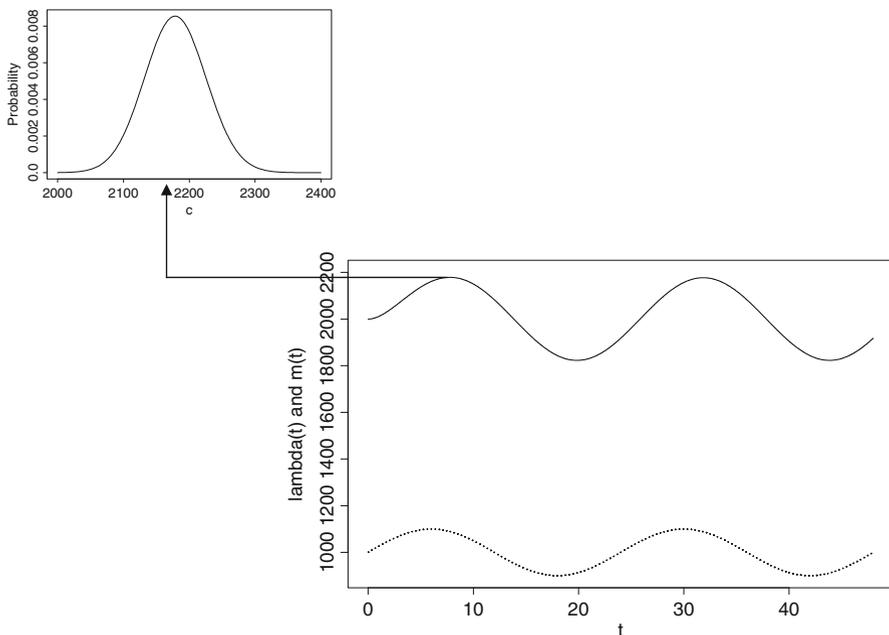


Fig. 3.1 Plot of $\lambda(t) = 1,000 + 100\sin(\pi t/12)$ (bottom) and $m(t)$ (top) when $\tau = 2$, and the Poisson distribution of number of cars in the garage with mean m^* (top left)

are an infinite number of servers. In Example 3.1 the customers are cars arriving at a rate of $\lambda(t)$ per hour, a “service time” is the time spent occupying a space with mean τ hours, and we pretend that there are infinitely many spaces. A good reference for using queuing models to address the design of service systems like call centers or the parking lot is Whitt (2007).

Let $N(t)$ be the number of cars in the parking garage at time $t \geq 0$. Then it is known that $N(t)$ has a Poisson distribution with mean $m(t)$, where $m(t)$ solves the differential equation

$$\frac{d}{dt}m(t) = \lambda(t) - \frac{m(t)}{\tau} \quad (3.1)$$

with an appropriate initial condition, say, $m(0) = 0$ if the garage starts empty (see Nelson and Taaffe 2004). One way to set the capacity of the parking garage is to solve (or numerically integrate over time) Eq. (3.1) to find the largest value attained by $m(t)$, say m^* . Then at its most congested, the distribution of the number of cars in the garage is Poisson with mean m^* , which implies that

$$\Pr\{\text{number of cars in the garage} \leq c\} = \sum_{n=0}^c \frac{e^{-m^*} (m^*)^n}{n!}. \quad (3.2)$$

Thus, c can be adjusted until this probability is sufficiently close to 1 (e.g., 0.99). Figure 3.1 shows a plot of the arrival rate to the garage, $\lambda(t)$, the mean number of

cars in the garage, $m(t)$, and the probability distribution of the number of cars in the garage at time t^* when the mean number is m^* (*Note: The Poisson is a discrete distribution, but is plotted as a continuous curve here.*)

Although there is no need to simulate an $M(t)/M/\infty$ queue to design the garage, only slightly enriching the model would make it mathematically intractable. For instance, in a large garage it might be important to account for drivers who have difficulty finding a space, or to include drivers who occupy multiple spaces to protect their cars and employees who also park in the garage but stay 8 h. These changes and others could necessitate simulation; however, the behavior would be similar to the $M(t)/M/\infty$ results.

If we needed to solve the parking lot problem via simulation, what issues would arise that we did not see in the TTF problem?

- Like the TTF problem, the system state will include (at least) the number of cars currently in the garage, and the events will include car arrivals and departures, similar to the failure and repair events in the TTF simulation. However, a single “Next Departure” event will not be sufficient²; in fact, there will need to be one pending departure event for each car currently in the garage. Effective ways to program simulations that do not have a small, fixed number of events are covered in Chap. 4.
- To build the simulation we need the facility to model and simulate a time-varying, stochastic arrival process and random parking times. Chapter 2 hinted at how we generate random variates, but what is needed here is clearly another level of sophistication. Methods for fitting and simulating input processes are covered in Chap. 6.
- We need to track $N(t)$, the number of cars in the garage, over some relevant time horizon. Unlike the average system availability in the TTF problem, the average number of cars in the infinite-sized garage is not relevant unless we are satisfied with having many cars turned away. The maximum observed value of $N(t)$ might seem better, but in a simulation we could, by chance, observe a one-in-a-million event; do we want to size the garage to protect against something that rare? How can we estimate something like (3.2) using simulation? These issues are covered in Chap. 7.

3.2 $M/G/1$ Queue

Example 3.2 (Hospital reception). The receptionist in a medium-sized hospital helps direct entering patients and visitors to the relevant floor or wing of the building. Under discussion is replacing the human receptionist with an electronic kiosk with a touch screen that might be a bit slower, or more variable, for patients and visitors who are less comfortable interacting with a touch screen. The hospital

² In the special case of exponentially distributed parking times one departure event is adequate due to the memorylessness property of the distribution; however, this will not be the case generally.

management engineers want to evaluate how much additional delay that this might cause.

For a rough-cut analysis of the situation the management engineers could use hospital records to estimate an overall arrival rate of patients and visitors to the reception desk and collect data on people interacting with the kiosk from a trial study with the vendor. Pretending that the characteristics of any system are unchanging over time is always an approximation, but such approximations often lead to useful and mathematically tractable models. Such is the case for the $M/G/1$ queue, a single-server queueing system to which customers arrive according to a Poisson process with rate λ customers/time, and whose service times are i.i.d. random variables with mean τ and standard deviation σ time units. The $M/G/1$ model is mathematically tractable, and illustrates several important ideas for computer simulation.

Let A_1, A_2, \dots be a sequence of i.i.d. random variables with mean $1/\lambda$ representing the interarrival times between customers (where A_1 is the actual time of the first arrival). Similarly, let X_1, X_2, \dots be the i.i.d. service times of successive customers with mean τ and standard deviation σ . Then if Y_1, Y_2, \dots are the successive waiting times in queue (i.e., the time from customer arrival until service begins), a little thought reveals that

$$Y_i = \max\{0, Y_{i-1} + X_{i-1} - A_i\}, \quad i = 1, 2, \dots \quad (3.3)$$

where we need to define $Y_0 = X_0 = 0$ to make the recursion work. This is known as Lindley's equation (Gross et al. 2008, p. 14), which not only is a convenient way to simulate the $M/G/1$ waiting-time process, but also provides a number of insights about many simulation output processes³:

- The successive waiting times are dependent, since Y_i clearly depends on Y_{i-1} (if the customer in front of me waited a long time, I probably will too, unless I arrive well after her).
- The outputs are not identically distributed. Clearly $Y_1 = 0$ (the first customer arrives to an empty system and does not wait), but the other Y_i 's are not certain to be 0.
- No matter how many customers we simulate ($i \rightarrow \infty$), the waiting times remain random variables.

This sounds like a difficult system to analyze, since the data are neither independent nor identically distributed, and they do not converge to some constant that we might use to summarize the system. What we might hope, however, is that there is a more complicated, but still useful, limiting behavior.

³ Lindley's equation makes queueing simulation look deceptively easy. However, if there are ten servers instead of one, or if customers are served in some priority order, or if we are interested in the number of customers in the queue rather than the waiting time, then an event-based simulation is typically required. Exercise 4 asks you to develop such a simulation; see also Chap. 4.

Provided $\lambda < \infty$ (customers do not arrive infinitely fast), $\rho = \lambda \tau < 1$ (on average we can serve customers at least a little faster than they arrive), and $\sigma < \infty$, then a number of things happen as $i \rightarrow \infty$:

1. The outputs Y_1, Y_2, \dots converge in distribution⁴ to a random variable Y whose distribution is not a function of the customer number i . That is, even though customer waiting times are always random variables, the distribution of those random variables is no longer changing. This makes the distribution of Y a nice summary of long-run performance, but it is not clear how to simulate it.
2. The sample mean $\bar{Y}(m) = m^{-1} \sum_{i=1}^m Y_i$ converges with probability 1 to a constant μ which we can interpret as the long-run average waiting time. This is less descriptive than the distribution of Y , but is still a useful summary measure. More importantly, it is easy to see how to estimate μ via simulation: Take the average of a very large number of simulated waiting times generated by using (3.3).
3. The $E(Y)$ and μ are equal, and are given by the Pollaczek–Khinchine formula (Gross et al. 2008, Sect. 5.1.1)

$$\mu = E(Y) = \frac{\lambda(\sigma^2 + \tau^2)}{2(1 - \lambda\tau)}. \quad (3.4)$$

Using this formula the management engineers can evaluate the impact of an increased mean service time (τ) and increased variability (σ^2) without the need to simulate. Unfortunately, as in the parking lot example, relatively small enhancements of the model to make it more realistic also make it mathematically intractable. However, (3.4) still provides insight about how arrival rates, mean service times and service variability affect queueing performance: As $\rho = \lambda \tau$ approaches 1 the mean waiting time increases explosively, while it increases linearly in the variance of the service times.

Suppose we used the natural estimator $\bar{Y}(m)$ to estimate μ via simulation. Obviously we have to choose the number of patients and visitors to simulate $m < \infty$ because we have to stop the simulation to obtain our estimate. How many patients and visitors is enough? One way to answer that question is by picking m large enough so that the standard error of $\bar{Y}(m)$, which is $\sqrt{\text{Var}(\bar{Y}(m))}$, is small enough. The standard error can be interpreted as the “average error” in $\bar{Y}(m)$ and often appears in confidence intervals.

In Chap. 5 we will discuss the fact that for certain types of simulation output processes the limit

$$\gamma^2 = \lim_{m \rightarrow \infty} m \text{Var}(\bar{Y}(m)) \quad (3.5)$$

exists, even though the output data Y_1, Y_2, \dots are not i.i.d. (it is easy to show that this limit exists when the data are i.i.d.). We call γ^2 the *asymptotic variance*, and it suggests that for m large enough, $\sqrt{\text{Var}(\bar{Y}(m))} \approx \gamma/\sqrt{m}$.

In the special case of an $M/M/1$ queue (which means the service time is exponentially distributed) and with $\tau = 1$, it can be shown that

⁴ The various modes of convergence (in distribution, with probability 1, and others) are defined in Chap. 5.

$$\gamma^2 = \frac{\rho(2 + 5\rho - 4\rho^2 + \rho^3)}{(1 - \rho)^4} \approx \frac{4\rho}{(1 - \rho)^4}$$

(Whitt 1989), so that

$$\sqrt{\text{Var}(\bar{Y}(m))} \approx \frac{2\sqrt{\rho}/(1 - \rho)^2}{\sqrt{m}}. \quad (3.6)$$

Notice that how much simulation effort is required depends dramatically on how close the system is to its capacity (i.e., how close ρ is to 1): For instance, m needs to be about 1,000 times larger to obtain the same standard error at $\rho = 0.9$ as at $\rho = 0.5$ (verify this for yourself using Eq. (3.6)). This is a feature of simulations that involve queues or networks of queues. Determination of the run length m , and estimation of γ^2 , are topics of Chap. 8.

3.3 The AR(1) Surrogate Model

The examples in this chapter are simplified, and therefore tractable, versions of realistic models that we might actually simulate. The lone exception is the AR(1) model presented here.

AR(1) is an abbreviation for “autoregressive order-1”; it is a model that arises naturally in time series forecasting (e.g., Chatfield 2004). Here we use it as a stand-in, or surrogate, for output from a stochastic simulation experiment because it shares the properties of many simulation output processes, but is very readily analyzed.

The AR(1) model for the i th simulation output (within a single replication) is

$$Y_i = \mu + \varphi(Y_{i-1} - \mu) + X_i, \quad i = 1, 2, \dots \quad (3.7)$$

with the following conditions:

- X_1, X_2, \dots are i.i.d. random variables with mean 0 and finite variance σ^2 ; thus, the X 's represent the stochastic nature of the simulation output.
- $|\varphi| < 1$; thus, φ controls the strength of dependence among successive outputs, with $\varphi = 0$ corresponding to i.i.d. output, $0 < \varphi < 1$ corresponding to positive dependence, and $-1 < \varphi < 0$ corresponding to negative dependence.
- μ is a constant; thus, μ is the long-run average, which is often the performance measure of interest.
- The distribution of Y_0 is given (including the typical case where it is a constant) and independent of the X 's; thus, Y_0 represents the initial condition at the beginning of the simulation which clearly has an impact on the outputs that follow.

The similarity of the AR(1) model to Lindley's equation (3.3) for the M/G/1 queue,

$$Y_i = \max\{0, Y_{i-1} + X_{i-1} - A_i\}, \quad i = 1, 2, \dots$$

is apparent, but the absence of the max operator makes the AR(1) easier to analyze. In particular, we can use induction to show that

$$Y_i = \mu + \varphi^i(Y_0 - \mu) + \sum_{j=0}^{i-1} \varphi^j X_{i-j}. \quad (3.8)$$

From Eq. (3.8) we can derive several properties of AR(1) output immediately:

$$E(Y_i) = \mu + \varphi^i (E(Y_0) - \mu) \xrightarrow{i \rightarrow \infty} \mu \quad (3.9)$$

$$\text{Var}(Y_i) = \varphi^{2i} \text{Var}(Y_0) + \sigma^2 \sum_{j=0}^{i-1} \varphi^{2j} \xrightarrow{i \rightarrow \infty} \frac{\sigma^2}{1 - \varphi^2}. \quad (3.10)$$

It can also be shown that $\text{Corr}(Y_i, Y_{i+j}) \xrightarrow{i \rightarrow \infty} \varphi^j$, implying that the correlation between outputs is geometrically decreasing with the number of observations (called the “lag”) between them.

Now suppose that the X ’s are normally distributed, and Y_0 is a constant. Then since Y_i is the sum of normally distributed random variables by (3.8), it is also normally distributed. And since a normally distributed random variable is completely characterized by its mean and variance, the limits in (3.9) and (3.10) show that as i increases (the length of the replication becomes longer) the distribution of the AR(1) output becomes independent of i , and specifically converges to a random variable $Y \sim N(\mu, \sigma^2/(1 - \varphi^2))$. This illustrates the concept of convergence in distribution, introduced for the M/G/1 queue in Sect. 3.2.

As compared to Lindley’s equation, the AR(1) model is particularly useful for mathematically evaluating the features of simulation output that impact output analysis because the key factors that might be important can be individually controlled: variability through σ^2 , dependence through φ , and initialization through Y_0 .

3.4 A Stochastic Activity Network

This example is based on Burt and Garman (1971) and Henderson and Nelson (2006, Chap. 1, Sect. 2.2).

Example 3.3 (Construction project). A construction project consists of a large number of activities. Some of these can be completed in parallel (drywall can be ordered while the foundation is being poured), while some cannot commence until others are completed (the roof cannot be constructed until the framing is finished). Since the durations of the activities are not precisely predictable, the project planners would like to take into account this variability when bidding the project because there will be penalties for completing the project after the contract date.

Project planning problems such as Example 3.3 can sometimes be modeled as stochastic activity networks (SANs); a small instance—which will be used for illustration here—is shown in Fig. 3.2. The nodes (circles) represent project milestones, and the arcs (arrows) are activities. The project starts with all activities emanating

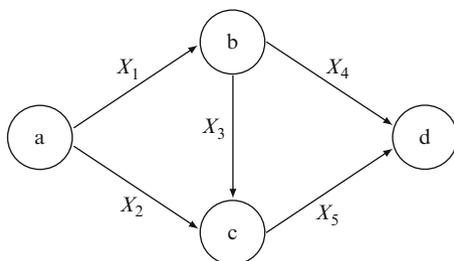


Fig. 3.2 A small stochastic activity network

from the source (first) node and is completed when the sink (last) node is reached. The rule is that all outgoing activities from a node commence when all of the incoming activities to that node are completed. The duration of the ℓ th activity is a random variable X_ℓ . Thus, the time to complete the project will be the longest path through the network:

$$Y = \max\{X_1 + X_4, X_1 + X_3 + X_5, X_2 + X_5\}. \quad (3.11)$$

The project planners are interested in information about the distribution of Y , for instance $\theta = \Pr\{Y > t_p\}$, where t_p is the quoted duration of the project.

Now if it happens that the activity durations are independent, and have common exponential distribution with mean 1, then a sequence of careful conditioning arguments gives (Burt and Garman 1971)

$$\Pr\{Y \leq t_p\} = \left(\frac{1}{2}t_p^2 - 3t_p - 3\right)e^{-2t_p} + \left(-\frac{1}{2}t_p^2 - 3t_p + 3\right)e^{-t_p} + 1 - e^{-3t_p}. \quad (3.12)$$

However, the problem quickly becomes intractable if some or all of the distributions are not exponential, or if X_1, X_2, \dots, X_5 are not independent. Then simulation is required. The following simulation could be used to estimate θ :

1. set $s = 0$
2. repeat n times:
 - a. generate X_1, X_2, \dots, X_5
 - b. set $Y = \max\{X_1 + X_4, X_1 + X_3 + X_5, X_2 + X_5\}$
 - c. if $Y > t_p$ then set $s = s + 1$
3. estimate θ by $\hat{\theta} = s/n$

Let Y_1, Y_2, \dots, Y_n be the n project completion times generated by this simulation. Assuming the X 's are generated independently on each trial, then the Y 's are naturally i.i.d. and classical statistical analysis applies.

Like some earlier examples, an event-based simulation is not required for this small problem; however, in a project with a realistic number of activities (perhaps

numbering in the hundreds) it may be prohibitively difficult to explicitly identify all of the paths through the network, but relatively easy to treat the completion of activities as events that may trigger the release of milestones. In this case an event-based simulation represents these paths implicitly, but is far easier to construct.

Let $\mathcal{I}(j)$ be the set of inbound activities to node j ; for instance, $\mathcal{I}(c) = \{2, 3\}$ for the example in Fig. 3.2. Similarly, let $\mathcal{O}(j)$ be the outgoing activities from node j ; thus, $\mathcal{O}(c) = \{5\}$. Finally, let $\mathcal{D}(\ell)$ be the destination node for activity ℓ , so that $\mathcal{D}(5) = d$. All of these sets would have to be known to construct the SAN, or could easily be extracted from a graphical representation. Now only a single type of event, which we call “milestone,” is required to simulate the SAN. The milestone event has as arguments a target node j and an activity ℓ that is inbound to that node. It removes the completed activity from the set of inbound activities to that node, $\mathcal{I}(j)$, and when that set becomes empty it schedules all of the node’s outbound activities $\mathcal{O}(j)$ to begin. In pseudocode, the event is as follows:

```

event milestone (activity  $\ell$  inbound to node  $j$ )
 $\mathcal{I}(j) = \mathcal{I}(j) - \ell$ 
if  $\mathcal{I}(j) = \emptyset$  then
    for each activity  $i \in \mathcal{O}(j)$ 
        schedule milestone(activity  $i$  inbound to node  $\mathcal{D}(i)$  to occur  $X_i$  time later)
end if

```

This representation makes it much easier to program a stochastic activity network simulation, but actually increases the computer effort required to simulate it due to the need to manage an event calendar and execute the events. Later in the book we will look for ways to make this simulation more statistically efficient, meaning that fewer replications will be needed to obtain the desired precision of the simulation estimate.

3.5 Asian Option

A “call” option is a contract giving the holder the right to purchase a stock for a fixed “strike price” at some time in the future. If the stock’s value increases well above the strike price, then the option is a good deal provided the purchase price for the option was not too high. If offered a call option with strike price K and maturity T on a stock that is currently trading at $X(0)$, how much should one be willing to pay for this option?

Valuing, pricing and creating financial instruments that serve various needs is the purview of the field of financial engineering (FE). FE often requires stochastic simulation. The critical issue in FE relative to the other examples in this chapter is the need for highly precise estimates. A 5% error may be tolerable in estimating customer delay in queue, but entirely out of the question in pricing a financial product that may be sold in millions of units. The development in this section is based on Glasserman (2004, Chap. 1), an excellent text on simulation and FE.

The value of the asset on which the option is written (e.g., stock) is often modeled as a continuous-time stochastic process $\{X(t), 0 \leq t \leq T\}$. In the standard European option the payoff to the contract holder is

$$(X(T) - K)^+ = \max\{0, X(T) - K\}.$$

That is, if the value of the stock at maturity, $X(T)$, is greater than the strike price K then the owner can exercise the option—purchase the stock at price K —and immediately sell it for a profit of $X(T) - K$. Of course, if $X(T) \leq K$, then the owner might as well just buy the stock on the open market and the option is worthless.

If the stochastic process $\{X(t), 0 \leq t \leq T\}$ is *geometric Brownian motion* (GBM, see Glasserman 2004, Sect. 3.2), then the value of a European option can easily be computed using the celebrated Black–Scholes formula (e.g., Glasserman 2004, Sect. 1.1). For the purpose of this book the important thing to know about GBM is that it is a continuous-time, continuous-state stochastic process; that is, $\{X(t), 0 \leq t \leq T\}$ has a value at all times t and $X(t) \in \mathfrak{R}$. This will cause some problems for simulation.

Since the European option can be valued without simulation, we take as our example a more difficult case:

Example 3.4 (Asian option). Suppose we are offered an “Asian” option, which means that the payoff is $(\bar{X}(T) - K)^+$, where $\bar{X}(T)$ is an average of the stock’s value over $0 \leq t \leq T$ (from now until maturity). Given the stock’s current value $X(0)$, a maturity T , a risk-free interest rate r (at which we could invest our money instead of buying the option), and a strike price K , the value of this option is

$$v = E \left[e^{-rT} (\bar{X}(T) - K)^+ \right]$$

where the multiplier e^{-rT} discounts the value at time T back to time 0 (now). How can we use simulation to evaluate it?

This would appear to be the easiest example introduced in this chapter, since the following algorithm does it:

1. set $s = 0$
2. repeat n times:
 - a. generate $\bar{X}(T)$
 - b. set $Y = e^{-rT} \max\{0, \bar{X}(T) - K\}$
 - c. set $s = s + Y$
3. estimate v by $\hat{v} = s/n$

The difficulty comes in Step 2a. We have not (yet) been precise about what is meant by $\bar{X}(T)$. Probably the most natural definition is

$$\bar{X}(T) = \frac{1}{T} \int_0^T X(t) dt. \tag{3.13}$$

We have seen time averages before (e.g., Sect. 1.2), but this one is different: $X(t)$ does not just change at discrete points in time, it changes continuously. Thus, if we wanted to try to schedule events at which the state of $X(t)$ changes, there would be an uncountably infinite number of them, even on a finite time interval $[0, T]$. This issue arises in many FE simulations because continuous-time, continuous-state stochastic processes are often used to represent the values of underlying financial assets, such as stocks, over time.

A natural approximation for (3.13) is to take the interval $[0, T]$, divide it up into m steps of size $\Delta t = T/m$ and use

$$\widehat{\bar{X}}(T) = \frac{1}{m} \sum_{i=1}^m X(i\Delta t)$$

the average value of the stock at a set of monitoring times $\Delta t, 2\Delta t, \dots, m\Delta t$.⁵ This introduces discretization error, of course, but not much if Δt is not too large. On the other hand, if Δt is too small then numerical round-off error can accumulate and the simulation can become quite slow. Thus, discrete approximations to continuous processes cannot be done without thought.

Suppose that we want to do a discrete approximation. If $X(t)$ is described by GBM with drift r (the risk-free interest rate) and volatility σ^2 (a measure of the variability of the asset's value), then it can be shown that for $0 = t_0 < t_1 < t_2 < \dots < t_m = T$,

$$X(t_{i+1}) = X(t_i) \exp \left\{ \left(r - \frac{1}{2} \sigma^2 \right) (t_{i+1} - t_i) + \sigma \sqrt{t_{i+1} - t_i} Z_{i+1} \right\} \quad (3.14)$$

where Z_1, Z_2, \dots, Z_m are i.i.d. $N(0, 1)$. We will use this fact to simulate the Asian option in Chap. 4.

Exercises

1. Show that the limit (3.5) exists, and precisely what it is, if Y_1, Y_2, \dots are i.i.d. with finite variance ζ^2 .
2. For the $M/M/1$ queue, make a plot that shows the relative effort required (value of m) to obtain equal standard error $\sqrt{\text{Var}(\bar{Y}(m))}$ for $0.5 \leq \rho \leq 0.99$. Note that for the $M/M/1$ queue Eq. (3.4) simplifies because $\sigma = \tau$ for the exponential distribution.
3. Instead of trying to obtain the same standard error $\sqrt{\text{Var}(\bar{Y}(m))}$ for $0.5 \leq \rho \leq 0.99$, we could instead try to obtain the same *relative error*

⁵ Many Asian options are actually defined in this way with the time steps specified as part of the contract. For the discussion here we will assume that the time average $\bar{X}(T)$ is what is actually desired, however.

$$\frac{\sqrt{\text{Var}(\bar{Y}(m))}}{E(Y)}$$

for $0.5 \leq \rho \leq 0.99$. Make a plot similar to Exercise 2 for relative error.

4. Develop an event-based simulation of the $M/G/1$ queue similar to the TTF example. Keep track of the number of customers in the system.
5. Derive (3.8) via induction.
6. Show how to represent the simulation of the SAN in Fig. 3.2 using an event-based approach.
7. Sketch out the event logic you would use to simulate the $M(t)/M/\infty$ queue to keep track of both the average number of cars in the parking lot and the maximum number in the lot for a 24-h period. Do not be concerned with how to simulate the car arrival process, just assume that you have a method for generating the interarrival times between cars.
8. Using (3.14), add details to the algorithm for simulating an Asian option in Sect. 3.5 to produce a discrete approximation. Let m , the number of steps, be an input.
9. For the SAN example in Sect. 3.4, how many replications are needed to estimate θ with relative error 1% if $t_p = 5$?

Chapter 4

Simulation Programming with VBASim

This chapter shows how simulations of some of the examples in Chap. 3 can be programmed in VBASim. The goals of the chapter are to introduce VBASim, and to hint at the experiment design and analysis issues that will be covered in later chapters. A complete listing of the VBASim source code can be found in Appendix A. This chapter can be skipped without loss of continuity; Java and Matlab versions of the source code and this chapter may be found at the book website.

4.1 VBASim Overview

VBASim is a collection of VBA subs, functions and class modules that aid in developing discrete-event simulations. They are entirely open source and can be modified to suit the user. The random-number and random-variate generation routines are VBA translations of the corresponding routines in simlib (Law 2007) which is written in C. VBASim is designed to be easy to understand and use, but not necessarily efficient.

Here is a brief description of the subs and functions in VBASim:

Sub `VBASimInit`: Initializes VBASim for use, typically called before the start of each replication.

Sub `Schedule`: Schedules future events.

Sub `SchedulePlus`: Schedules future events and allows an object to be stored with the event.

Sub `Report`: Writes a result to a specific row and column of an Excel worksheet.

Sub `ClearStats`: Clears certain statistics being recorded by VBASim.

Sub `InitializeRNSeed`: Initializes the random-number generator; typically called only once in a simulation.

Function `Expon`: Generates exponentially distributed random variates.

Function `Uniform`: Generates uniformly distributed random variates.

Function `Random_integer`: Generates a random integer.
 Function `Erlang`: Generates Erlang distributed random variates.
 Function `Normal`: Generates normally distributed random variates.
 Function `Lognormal`: Generates lognormally distributed random variates.
 Function `Triangular`: Generates triangularly distributed random variates.

The random-variate generation functions take two types of arguments: parameters, and a random-number stream; the random-number stream is always the last argument. For instance

```
X = Uniform(10, 45, 2)
```

generates random-variates that are uniformly distributed between 10 and 45, using stream 2.

As you already know, the pseudorandom numbers we use to generate random variates in simulations are essentially a long list. Random-number streams are just different starting places in this list, spaced far apart. The generator in VBASim (which is a translation of the generator in Law 2007) has 100 streams. Calling `Sub InitializeRNSeed` sets the initial position of each stream. Then each subsequent call to a variate-generation routine using stream # advances stream # to the next pseudorandom number. The need for streams is discussed in Chap. 7, but it is important to know that any stream is as good as any other stream in a well-tested generator.

Here is a brief description of the class modules (objects) in VBASim:

`Entity`: Object used to model transient items that flow through the system.
`FIFOQueue`: Object to hold entities in first-in-first-out order.
`Resource`: Object used to model a scarce quantity.
`EventNotice`: Object used to represent a future event.
`EventCalendar`: Data structure holding event notices in chronological order.
`CTStat`: Object used to record continuous-time statistics.
`DTStat`: Object used to record discrete-time statistics.

4.2 Simulating the $M(t)/M/\infty$ Queue

Here we consider the parking lot example of Sect. 3.1, a queueing system with time-varying car arrival rate, exponentially distributed parking time and an infinite number of parking spaces. The simulation program consists of some global declarations (Fig. 4.1), a main program and some event routines (Fig. 4.2), an initialization sub (Fig. 4.3), a function to generate car arrivals (Fig. 4.4), and the support functionality provided by VBASim. Two important aspects of VBASim are illustrated by this model: event scheduling and collecting time-averaged statistics.

The key state variable that we need to track is the number of cars in the lot. There are two essential events that affect the state, the arrival of a car and the departure of a car, and a third event we will use to stop the simulation at a specific point in

time. The tricky part is that there can be an unlimited number of pending “departure” events; in fact, there are as many pending departure events as there are cars in the lot. Therefore, having a unique variable to represent the scheduled time of each pending event, as was used for the TTF example in Sect. 2, will not work.

To handle event scheduling, VBASim has an event calendar data structure called `Calendar`, a sub called `Schedule` for putting events on the `Calendar`, and a method called `Remove` for extracting the next event from the calendar. An event in VBASim is a VBA object of class `EventNotice` which has (at least) two properties: `EventType` and `EventTime`. The statement

```
Schedule (Name, Increment)
```

creates an `EventNotice`, assigns the character string `Name` to the `EventType` property, assigns the value `Clock + Increment` to the `EventTime` property, and schedules the `EventNotice` on the `Calendar` in the correct chronological order. The method `Calendar.Remove` extracts the chronologically next `EventNotice` from the `Calendar`, making its `EventType` and `EventTime` properties available to advance time and execute the proper event.

The simulation main program `Sub MtMInf` in Fig. 4.2 illustrates how the event-related features of VBASim work. The following four statements, or ones very similar, will appear in all simulations using VBASim:

```
Dim NextEvent As EventNotice

Set NextEvent = Calendar.Remove
Clock = NextEvent.EventTime
Select Case NextEvent.EventType
```

Since our simulation will repeatedly remove the next `EventNotice` from `Calendar`, we need an object of this class to which to assign it; the statement `Dim NextEvent As EventNotice` provides this. The statement `Set NextEvent = Calendar.Remove` illustrates how the `Remove` method extracts the next event, after which we advance the simulation `Clock` to time `NextEvent.EventTime` and execute the event indicated by `NextEvent.EventType`.

VBASim’s `Schedule` sub places events on the `Calendar`; for instance, the statement

```
Call Schedule("EndSimulation", 24)
```

creates an `EventNotice` of type `EndSimulation` to occur 24h from the time currently on `Clock` (which is 0 at the beginning of the simulation). Notice that VBASim requires the user to decide on the base time unit in the simulation and to use it consistently throughout. In this simulation the base time unit is hours.

The key state variable in this simulation is `N`, the current number of cars in the parking lot, and we want statistics on it. VBASim contains a `CTStat` class that can be used to record time-averaged statistics. Here is how it works: First, we declare a new `CTStat` using the statement

```

' Example illustrating use of VBASim for simulation
' of M(t)/M/infinity Queue parking lot example.
' In this version parking time averages 2 hours;
' the arrival rate varies around 100 per hour;
' the lot starts empty, and we look at a 24-hour period.
'
' See VBASim module for generic declarations
' See Class Modules for the supporting VBASim classes
' Parameters we may want to change
Dim MeanParkingTime As Double

' Simulation variables and statistics
Private N As Integer           'Number in queue
Dim QueueLength As New CTStat 'Use to keep statistics on N
Private MaxQueue As Integer    'Largest observed value of N

```

Fig. 4.1 Declarations for the parking lot simulation

```
Dim QueueLength As New CTStat
```

as shown in the global declarations (Fig. 4.1). Second, the `CTStat` can (and usually should) be added to a special collection called `TheCTStats`, as shown in Fig. 4.3. `VBASim` reinitializes any `CTStat` in `TheCTStats` collection whenever `Call VBASimInit` is executed, which is typically at the beginning of each replication. Next, whenever the value of the variable of interest changes, the `Record` method of the `CTStat` is employed to record the change (which means it is called just after the change occurs); in this simulation the statement is

```
QueueLength.Record (N)
```

as shown in Fig. 4.2. Finally, the time-averaged value of the `CTStat` can be computed using the `Mean` method of the `CTStat`, as in `QueueLength.Mean`.

Notice that `VBASim` also has a `Report` sub that writes a value or character string to a given row and column of an Excel worksheet. The syntax is

```
Call Report(value or string, worksheet, row, column)
```

See Figs. 4.2 and 4.3.

Recall that the arrival rate (in cars per hour) to the parking lot was modeled by the function $\lambda(t) = 1,000 + 100\sin(\pi t/12)$. To make the simulation execute more quickly for the purpose of this introduction, we change that rate to $\lambda(t) = 100 + 10\sin(\pi t/12)$, so that the arrival rate varies between 90 and 110 cars per hour, depending on the hour of the day t .

The Function `NSPP` shown in Fig. 4.4 generates interarrival times (time gaps between arrivals) from a nonstationary Poisson process with this arrival rate. The formal definition of a nonstationary Poisson process is a topic of Chap. 6. However, we provide an intuitive justification for how Function `NSPP` works here:

```

Private Sub MtMInf ()
' Initialize
  Dim Reps As Integer
  Dim NextEvent As EventNotice
  Call MyInit ' special one-time initializations
  For Reps = 1 To 1000
    ' Initializations for each replication
    N = 0
    MaxQueue = 0
    Call VBASimInit ' initialize VBASim for each replication
    Call Schedule("Arrival", NSPP(1))
    Call Schedule("EndSimulation", 24)
    Do
      Set NextEvent = Calendar.Remove
      Clock = NextEvent.EventTime
      Select Case NextEvent.EventType
        Case "Arrival"
          Call Arrival
        Case "Departure"
          Call Departure
      End Select
      Loop Until NextEvent.EventType = "EndSimulation"
    ' Write output report for each replication
    Call Report(QueueLength.Mean, "MtMInf", Reps + 1, 1)
    Call Report(MaxQueue, "MtMInf", Reps + 1, 2)
  Next Reps
  End ' ends execution, closes files, etc.
End Sub

Private Sub Arrival()
' Arrival event
' Schedule next arrival
  Call Schedule("Arrival", NSPP(1))
' Update number in queue and max
  N = N + 1
  QueueLength.Record (N)
  If N > MaxQueue Then
    MaxQueue = N
  End If
' Schedule departure
  Call Schedule("Departure", Expon(MeanParkingTime, 2))
End Sub

Private Sub Departure()
' End of service event
' Update number in queue
  N = N - 1
  QueueLength.Record (N)
End Sub

```

Fig. 4.2 Main program and event routines for the parking lot simulation

```

Private Sub MyInit ()

' Initialize the simulation
  Call InitializeRNSeed
  MeanParkingTime = 2
  TheCTStats.Add QueueLength

' Write headings for the output reports

  Call Report("Average Number In Queue", "MtMInf", 1, 1)
  Call Report("Maximum Number in Queue", "MtMInf", 1, 2)

End Sub

```

Fig. 4.3 Initializing the parking lot simulation

```

Private Function NSPP(Stream As Integer) As Double
' This function implements thinning to generate interarrival times
' from the nonstationary Poisson arrival process representing
' car arrivals. Time units are hours.

Dim PossibleArrival As Double

PossibleArrival = Clock + Expon(1 / 110, Stream)

Do Until Uniform(0, 1, Stream) < _
  (100 + 10 * VBA.Sin(3.141593 * PossibleArrival / 12)) / 110
  PossibleArrival = PossibleArrival + Expon(1 / 110, Stream)
Loop

NSPP = PossibleArrival - Clock

End Function

```

Fig. 4.4 Function to generate interarrival times to the parking lot

A stationary Poisson process has times between arrivals that are exponentially distributed with a fixed rate λ (or equivalently a constant mean time between arrivals $1/\lambda$). The inverse cdf method for generating exponentially distributed random variates was described in Sect. 2.2.1. The maximum arrival rate for $\lambda(t)$ is 110 cars per hour, so Function NSPP generates possible arrivals using a stationary arrival process with rate $\lambda = 110$. To achieve the time-varying arrival rate, it only accepts a *possible* arrival at time t as an *actual* arrival with probability $\lambda(t)/\lambda$. Thus, a possible arrival that is to occur at a time when $\lambda(t) = 110$ will always be an actual arrival, while a possible arrival that is to occur when $\lambda(t) = 90$ will only have a 90/110 probability of being an actual arrival. That this method, which is called “thinning,” generates a nonstationary Poisson process with the desired rate is discussed in Chap. 6.

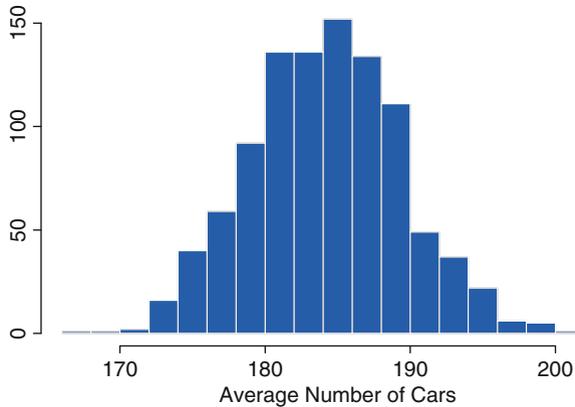


Fig. 4.5 Histogram of the daily average number of cars

Figure 4.5 shows a histogram of the 1,000 daily averages of the number of cars in the parking lot obtained by running the simulation for 1,000 replications; the overall average of these averages is 184.2 ± 0.3 cars, where the “ ± 0.3 ” part comes from a 95% confidence interval on the mean (confidence intervals are a subject of Chap. 7). Thus, the simulation provides a pretty precise estimate of the time-average mean number of cars that would be found in the (conceptually) infinite-size garage during a day.

The histogram shows that the average number of cars in the garage can vary substantially from day to day, so we certainly would not want to build a garage with a capacity of, say, 185 cars. Further, averaging over the day masks the largest number of cars in the garage during the day, and that number is more useful for selecting a finite capacity for the garage.

VBASim provides no special support for the maximum statistic, but since we have access to everything in a VBASim simulation we can easily record whatever we want. Here we define a variable `MaxQueue` which is initialized to 0 at the beginning of each replication, and is increased to the current value of `N` whenever `N` exceeds the previous value of `MaxQueue`. See in particular `Sub Arrival` in Fig. 4.2.

Suppose that we want the parking lot to be of adequate size 99% of the time. Since we record the maximum size on 1,000 replications, we could use the 990th value (sorted from smallest to largest) as the size of the garage, which turned out to be 263 cars in this simulation. Figure 4.6 shows the empirical cumulative distribution (ecdf) of the 1,000 maximums recorded by the simulation. The ecdf treats each observed value as equally likely (and therefore as having probability $1/1,000$), and plots the sorted maximum values on the horizontal axis and the cumulative probability of each observation on the vertical axis. The plot shows how the cumulative probability of 0.99 maps to the value of 263 cars, which might be a very reasonable capacity for the garage. Putting a confidence interval on this value is quite different than putting one on the mean, and will be discussed in Chap. 7. Without a confidence interval (or some measure of error) we cannot be sure if 1,000 replications is really enough to estimate the 99th percentile of the maximum.

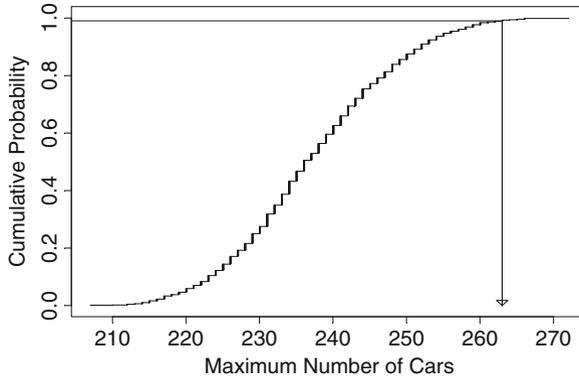


Fig. 4.6 Empirical cdf of the daily maximum number of cars in the parking lot

4.2.1 Issues and Extensions

1. The $M(t)/M/\infty$ simulation presented here simulates 24 h of parking lot operation, and treats each 24-h period as an independent replication starting with an empty garage. This only makes sense if the garage is emptied each day, for instance, if the mall closes at night. Is the assumed arrival rate $\lambda(t)$ appropriate for a mall that closes at night?
2. Suppose that the parking lot serves a facility that is actually in operation 24 h a day, 7 days per week (i.e., all the time). How should the simulation be initialized, and how long should the run length be in this case?
3. How could the simulation be initialized so that there are 100 cars in the parking lot at time 0?
4. When this example was introduced in Sect. 3.1, it was suggested that we size the garage based on the (Poisson) distribution of the number of cars in the garage at the point in time when the mean number in the garage was maximum. Is that what we did, empirically, here? If not, how is the quantity we estimated by simulation related to the suggestion in Sect. 3.1 (for instance, will the simulation tend to suggest a bigger or smaller garage than the analytical solution in Sect. 3.1)?
5. One reason that this simulation executes quite slowly when $\lambda(t) = 1,000 + 100\sin(\pi t/12)$ is that the thinning method we used is very inefficient (lots of possible arrivals are rejected). Speculate about ways to make it faster.
6. For stochastic processes experts: Another reason that the simulation is slow when $\lambda(t) = 1,000 + 100\sin(\pi t/12)$ is that there can be 1,000 or more pending departure events on Calendar at any time, which means that scheduling a new event in chronological order involves a slow search. However, it is possible to exploit the memoryless property of the exponential distribution of parking time to create an equivalent simulation that has only two pending events (the next car arrival and next car departure) at any point in time. Describe how to do this.

4.3 Simulating the $M/G/1$ Queue

Here we consider the hospital example of Sect. 3.2, a queueing system with Poisson arrival process, some (as yet unspecified) service-time distribution, and a single server (either a receptionist or an electronic kiosk); in other words, an $M/G/1$ queue. Patient waiting time is the key system performance measure, and the long-run average waiting time in particular.

Recall that Lindley's equation (3.3) provides a shortcut way to simulate successive customer waiting times:

$$\begin{aligned} Y_0 &= 0 & X_0 &= 0 \\ Y_i &= \max\{0, Y_{i-1} + X_{i-1} - A_i\}, & i &= 1, 2, \dots \end{aligned}$$

where Y_i is the i th customer's waiting time, X_i is that customer's service time, and A_i is the interarrival time between customers $i - 1$ and i . Lindley's equation avoids the need for an event-based simulation, but is limited in what it produces (how would you track the time-average number of customers in the queue?). In this section we will describe both recursion-based and event-based simulations of this queue, starting with the recursion.

4.3.1 Lindley Simulation of the $M/G/1$ Queue

To be specific, suppose that the mean time between arrivals is 1 min, with the distribution being exponential, and the mean time to use the kiosk is 0.8 min (48 s), with the distribution being an Erlang-3. An Erlang- p is the sum of p i.i.d. exponentially distributed random variables, so an Erlang-3 with mean 0.8 is the sum of three exponentially distributed random variables each with mean 0.8/3.

In Sect. 3.2 we noted that the waiting-time random variables Y_1, Y_2, \dots converge in distribution to a random variable Y , and it is $\mu = E(Y)$ that we will use to summarize the performance of the queueing system. We also noted that $\bar{Y}(m) = m^{-1} \sum_{i=1}^m Y_i$ converges with probability 1 to μ as the number of customers simulated m goes to infinity.

All of this suggests that we make a very long simulation run (large m) and estimate μ by the average of the observed waiting times Y_1, Y_2, \dots, Y_m . But this is not what we will do, and here is why: Any m we pick is not ∞ , so the waiting times early in the run—which will tend to be smaller than μ because the queue starts empty—will likely pull $\bar{Y}(m)$ down. To reduce this effect, we will let the simulation generate waiting times for a while (say d of them) before starting to actually include them in our average. We will still make m large, but our average will only include the last $m - d$ waiting times. That is, we will use as our estimator the truncated average

$$\bar{Y}(m, d) = \frac{1}{m - d} \sum_{i=d+1}^m Y_i. \quad (4.1)$$

```

Private Sub Lindley()

Dim Rep As Integer
Dim m As Long, i As Long
Dim d As Integer
Dim Y As Double, X As Double, A As Double
Dim SumY As Double

Call Report("Average Wait", "Lindley", 1, 1)
Call InitializeRNSeed
m = 55000
d = 5000

For Rep = 1 To 10
  Y = 0
  SumY = 0
  For i = 1 To d
    A = Expon(1, 1)
    X = Erlang(3, 0.8, 2)
    Y = WorksheetFunction.Max(0, Y + X - A)
  Next i
  For i = d + 1 To m
    A = Expon(1, 1)
    X = Erlang(3, 0.8, 2)
    Y = WorksheetFunction.Max(0, Y + X - A)
    SumY = SumY + Y
  Next i
  Call Report(SumY / CDBl(m - d), "Lindley", Rep + 1, 1)
Next Rep
End

End Sub

```

Fig. 4.7 Simulation of the $M/G/1$ queue using Lindley's equation

In addition, we will not make a single run of m customers, but instead will make n replications. This yields n i.i.d. averages $\bar{Y}_1(m, d), \bar{Y}_2(m, d), \dots, \bar{Y}_n(m, d)$ to which we can apply standard statistical analysis. This avoids the need to directly estimate the asymptotic variance γ^2 , a topic we defer to later chapters.

Figure 4.7 shows a VBASim simulation of the $M/G/1$ queue using Lindley's equation. In this simulation $m = 55,000$ customers, we discard the first $d = 5,000$ of them, and make $n = 10$ replications. The ten replication averages are individually written to an Excel worksheet called "Lindley" and are displayed in Table 4.1.

Notice that the average waiting time is a bit over 2 min, and that VBA, like all programming languages, displays a very large number of output digits. How many are really meaningful? A confidence interval is one way to provide an answer.

Since the across-replication averages are i.i.d., and since each across-replication average is itself the within-replication average of a large number of individual waiting times (50,000 to be exact), the assumption of independent, normally distributed

Table 4.1 Ten replications of the $M/G/1$ queue using Lindley's equation

Replication	$\bar{Y}(55,000, 5,000)$
1	2.191902442
2	2.291913404
3	2.147858324
4	2.114346960
5	2.031447995
6	2.110924602
7	2.132711743
8	2.180662859
9	2.139610760
10	2.146212039
Average	2.148759113
std. dev.	0.066748617

output data is reasonable. This justifies a t -distribution confidence interval on μ . The key ingredient is $t_{1-\alpha/2, n-1}$, the $1 - \alpha/2$ quantile of the t distribution with $n - 1$ degrees of freedom. If we want a 95% confidence interval, then $1 - \alpha/2 = 0.975$, and our degrees of freedom are $10 - 1 = 9$. Since $t_{0.975, 9} = 2.26$, we get $2.148759113 \pm (2.26)(0.066748617)/\sqrt{10}$ or $2.148759113 \pm 0.047703552$. This implies that we can claim with high confidence that μ is around 2.1 min, or we could give a little more complete information as 2.14 ± 0.05 min. Any additional digits are statistically meaningless.

Is an average of 2 min too long to wait? To actually answer that question would require some estimate of the corresponding wait to see the receptionist, either from observational data or a simulation of the current system. Statistical comparison of alternatives is a topic of Chap. 8.

4.3.2 Event-Based Simulation of the $M/G/1$ Queue

The simulation program consists of some global declarations (Fig. 4.8), a main program (Fig. 4.9), some event routines (Fig. 4.10), an initialization sub (Fig. 4.11), and the support functionality provided by VBASim. Four VBASim class objects and one sub are illustrated by this model: Entity, FIFOQueue, Resource, DTStat and ClearStats. At a high level, here is what they do:

- Entity objects are used to model transient items, such as transactions or customers that pass through the system. They can have attributes (VBA properties) that they carry with them; by default they have an attribute called CreateTime which is set to the value of Clock when an Entity object is created. In this simulation the Entity objects represent patients or visitors.
- FIFOQueue is a VBA collection, much like Calendar, that is used for holding Entity objects in first-in-first-out order; it also records time-average number in

```

' Example illustrating use of VBASim for
' simulation of M/G/1 Queue.

' See VBASim module for generic declarations
' See Class Modules for the supporting VBASim classes

' Parameters we may want to change
Private MeanTBA As Double      ' mean time between arrivals
Private MeanST As Double      ' mean service time
Private Phases As Integer     ' number of phases in service
                               ' distribution
Private RunLength As Double   ' run length
Private WarmUp As Double     ' "warm-up" time

' Global objects needed for simulation
' These will usually be queues and statistics

Dim Queue As New FIFOQueue    'customer queue
Dim Wait As New DTStat        'discrete-time statistics
                               'on customer waiting
Dim Server As New Resource    'server resource

```

Fig. 4.8 Declarations for the hospital simulation

the queue statistics. In this simulation the queue represents the patients using or waiting for the kiosk.

- Resource objects represent scarce quantities such as workers, machines and computers that are needed (typically) to serve or process an Entity in some way; they also record the average number of busy resource units. In this simulation the Resource object represents the kiosk.
- DTStat is an object for recording discrete-time statistics; it is the companion to CTStat for continuous-time statistics. In this simulation a DTStat is used to record total waiting-time statistics, where we will define “total waiting time” to be the time from patient arrival until they are finished with the kiosk. Notice that this is (intentionally) different from the definition in Sect. 4.3.1, where the wait only included the time to reach the front of the line.
- ClearStats is a sub that reinitializes all statistical variables found in two collections, TheCTStats and TheDTStats. FIFOQueue and Resource objects each create a CTStat which is automatically added to TheCTStats. When the programmer creates a custom CTStat or DTStat then they must add it to the appropriate collection.

Figure 4.8 shows the declarations of the FIFOQueue, DTStat and Resource objects, specifically:

```

Dim Queue As New FIFOQueue
Dim Wait As New DTStat
Dim Server As New Resource

```

```

Private Sub MG1()
' Initialize

    Dim Reps As Integer
    Dim NextEvent As EventNotice

    Call MyInit ' special initializations for this simulation

    For Reps = 1 To 10

        Call VBASimInit 'initialize VBASim for each replication

        Call Schedule("Arrival", Expon(MeanTBA, 1))
        Call Schedule("EndSimulation", RunLength)
        Call Schedule("ClearIt", WarmUp)

        Do
            Set NextEvent = Calendar.Remove
            Clock = NextEvent.EventTime
            Select Case NextEvent.EventType
            Case "Arrival"
                Call Arrival
            Case "EndOfService"
                Call EndOfService
            Case "ClearIt"
                Call ClearStats
            End Select
            Loop Until NextEvent.EventType = "EndSimulation"

        ' Write output report for each replication

        Call Report(Wait.Mean, "MG1", Reps + 1, 1)
        Call Report(Queue.Mean, "MG1", Reps + 1, 2)
        Call Report(Queue.NumQueue, "MG1", Reps + 1, 3)
        Call Report(Server.Mean, "MG1", Reps + 1, 4)

    Next Reps

    End ' ends execution, closes files, etc.

End Sub

```

Fig. 4.9 Main program for the hospital simulation

These are declared globally because there is only one of each and they will be referenced from many places in the simulation code.

This contrasts with the Entity objects that will be created and discarded as needed to represent patients coming and going. For instance, consider these three statements in Sub Arrival:

```

Dim Customer As New Entity
Queue.Add Customer
Set Customer = Nothing

```

The keyword `New` in the `Dim` statement causes a new instance of the `Entity` class to be created. We can access its attributes using the `.` extension, as in `Customer.CreateTime`. The `Add` method of the `FIFOQueue` object `Queue` places `Entity` objects into the queue in order of arrival. Finally, the `Set Customer = Nothing` statement frees the name `Customer` to be used again for a new `Entity`. This statement is not absolutely necessary unless we execute the statement `Dim Customer As New Entity` a second time within `Sub Arrival` because the name `Customer` will automatically be released when we exit the `Sub`. However, we think it is good programming practice to release the object name when it is no longer needed so there is no confusion. Notice that the `Entity` object is not lost because it is stored in `Queue`.

Moving into the `Sub EndOfService` event routine, the following four statements remove the first customer from the queue, use its `CreateTime` attribute to compute the total waiting time, record this value using the `DTStat` object `Wait`, and finally release the name `DepartingCustomer`. It is important to notice the *absence* of the keyword `New` in `Dim DepartingCustomer As Entity`; this means that `DepartingCustomer` is declared to be of type `Entity`, but a new `Entity` object is not created by the `Dim` statement.

```
Dim DepartingCustomer As Entity
Set DepartingCustomer = Queue.Remove
Wait.Record (Clock - DepartingCustomer.CreateTime)
Set DepartingCustomer = Nothing
```

Before a `Resource` object can be used, its capacity (number of identical units) must be set. In `Sub MyInit` this is accomplished by using the object's `SetUnit` method, `Server.SetUnits (1)`. If, for instance, there were three identical kiosks then this statement would be `Server.SetUnits (3)`. To make one (or more) units of a `Resource` busy, the `Seize` method is employed, while idling a `Resource` is accomplished via the `Free` method, as shown in the event subs.

In this simulation we are interested in long-run performance, so the length of a replication is determined by whatever we decide is long enough (which is not an easy decision, actually). When we used Lindley's equation it was natural to specify the replication length in terms of the number of customers simulated. However, in more complex simulations with many different outputs, it is far more common to specify the replication length by a stopping time T chosen so that it will be long enough for all of the outputs. Similarly, if we plan to discard data, then it is easier to specify a time T_d at which all statistics will be cleared.

To control the replication length by time, we schedule two events: an "EndSimulation" event scheduled to occur at time $T = 55,000$ min, and a statistics clearing event "ClearIt" that calls the sub `ClearStats` at time $T_d = 5,000$ min. *Because the arrival rate of patients and visitors is 1 per minute, these times will approximately correspond to the run lengths of 55,000 and 5,000 patients; however the actual number of patients will be random and vary from replication to replication.*

```

Private Sub Arrival()
' Arrival event

' Schedule next arrival
  Call Schedule("Arrival", Expon(MeanTBA, 1))

' Process the newly arriving customer

  Dim Customer As New Entity
  Queue.Add Customer
  Set Customer = Nothing
' If server is not busy, start service by seizing the server

  If Server.Busy = 0 Then
    Server.Seize (1)
    Call Schedule("EndOfService", Erlang(Phases, MeanST, 2))
  End If

End Sub

Private Sub EndOfService()
' End of service event

' Remove departing customer from queue and record wait time

  Dim DepartingCustomer As Entity
  Set DepartingCustomer = Queue.Remove
  Wait.Record (Clock - DepartingCustomer.CreateTime)
  Set DepartingCustomer = Nothing

' Check to see if there is another customer; if yes start service
' otherwise free the server

  If Queue.NumQueue > 0 Then
    Call Schedule("EndOfService", Erlang(Phases, MeanST, 2))
  Else
    Server.Free (1)
  End If

End Sub

```

Fig. 4.10 Event routines for the hospital simulation

Let $\bar{Y}(T, T_d)$ be a replication average of all waiting times recorded between times T_d and T . Clearly $\bar{Y}(m, d)$ based on count, and $\bar{Y}(T, T_d)$ based on time, will have different statistical properties, but it is intuitively clear that both will be good estimators of μ if their arguments are fixed (not a function of the data) and large enough. Notice also that a run time and deletion time are ideal for continuous-time statistics like the time-average number in queue.

```

Private Sub MyInit ()

' Initialize the simulation
  Call InitializeRNSeed
  Server.SetUnits (1) ' set the number of servers to 1
  MeanTBA = 1
  MeanST = 0.8
  Phases = 3
  RunLength = 55000
  WarmUp = 5000

' Add queues, resources and statistics that need to be
' initialized between replications to the global collections

  TheDTStats.Add Wait
  TheQueues.Add Queue
  TheResources.Add Server

' Write headings for the output reports

  Call Report("Average Wait", "MG1", 1, 1)
  Call Report("Average Number in Queue", "MG1", 1, 2)
  Call Report("Number Remaining in Queue", "MG1", 1, 3)
  Call Report("Server Utilization", "MG1", 1, 4)

End Sub

```

Fig. 4.11 Initializing the hospital simulation

For this event-based simulation it is easy to record and report a number of statistics. `FIFOQueue`, `Resource` and `DTStat` objects all have `Mean` methods for reporting average values, which for this simulation are `Queue.Mean`, `Server.Mean` and `Wait.Mean`, respectively. In addition, the `FIFOQueue` objects have a `NumQueue` method to deliver the current number of entities in the queue; here we use it to report the number of patients in the queue at time 55,000 when the replication ends. Table 4.2 shows the results from ten replications, along with the overall averages and 95% confidence interval half-widths.

Again there are meaningless digits, but the confidence intervals can be used to prune them. For instance, for the mean total wait we could report 2.93 ± 0.05 min. How does this statistic relate to the 2.14 ± 0.05 min reported for the simulation via Lindley's equation? Mean total time in the kiosk system (which is what the event-based simulation estimates) consists of mean time waiting to be served (which is what the Lindley simulation estimates) plus the mean service time (which we know to be 0.8 min). So it is not surprising that these two estimates differ by about 0.8 min.

Table 4.2 Ten replications of the $M/G/1$ queue using the event-based simulation

Rep	Total wait	Queue	Remaining	Utilization
1	2.996682542	3.017423450	1	0.806654823
2	3.103155842	3.127773149	7	0.807276539
3	2.951068607	2.948352414	2	0.799341091
4	2.848547497	2.846673097	0	0.794701908
5	2.908913572	2.900437432	2	0.798615617
6	2.895622648	2.896900635	3	0.801983001
7	2.909777649	2.901219722	0	0.798658678
8	2.914666297	2.908612119	4	0.795440658
9	2.922193762	2.923535588	0	0.799157017
10	2.873148311	2.862172885	0	0.799143690
Average	2.932377673	2.933310049	1.9	0.800097302
Std. dev.	0.07227642	0.082882288	2.282785822	0.004156967
$\pm 95\%$ CI	0.051654132	0.059233879	1.631449390	0.002970879

4.3.3 Issues and Extensions

1. In what situations does it make more sense to compare the simulated kiosk system to simulated data from the current receptionist system rather than real data from the current receptionist system?
2. It is clear that if all we are interested in is mean waiting time, defined either as time until service begins or total time including service, the Lindley approach is superior (since it is clearly faster, and we can always add in the mean service time to the Lindley estimate). However, if we are interested in the distribution of total waiting time, then adding in the mean service time does not work. How can the Lindley recursion be modified to simulate total waiting times?
3. How can the event-based simulation be modified so that it also records waiting time until service begins?
4. How can the event-based simulation be modified to clear statistics after exactly 5,000 patients, and to stop at exactly 55,000 patients?
5. The experiment design method illustrated in the event-based simulation is often called the “replication–deletion” method. If we only had time to generate 500,000 waiting times, what issues should be considered in deciding the values of n (replications), m (run length) and d (deletion amount)? Notice that we must have $nm = 500,000$, and only $n(m - d)$ observations will be used for estimating the mean μ .
6. An argument against summarizing system performance by long-run measures is that no system stays unchanged forever (55,000 patients is approximately 38 24-h days during which time there could be staff changes, construction or emergencies), so a measure like μ is not a reflection of reality. The counter-argument is that it is difficult, if not impossible, to model all of the detailed changes that occur over any time horizon (even the time-dependent arrival process in the $M(t)/M/\infty$ simulation is difficult to estimate in practice), so long-run performance at least provides an understandable summary measure (“If our process stayed the same,

then over the long run. . .”). Also, it is often mathematically easier to obtain long-run measures than it is to estimate them by simulation (since simulations have to stop). Considering these issues, what sort of analysis makes the most sense for the hospital problem?

4.4 Simulating the Stochastic Activity Network

Here we consider the construction example of Sect. 3.4 which is represented as a stochastic activity network (SAN). Recall that the time to complete the project, Y , can be represented as

$$Y = \max\{X_1 + X_4, X_1 + X_3 + X_5, X_2 + X_5\}$$

where X_i is the duration of the i th activity. This simple representation requires that we enumerate all paths through the SAN, so that the project completion time is the longest of these paths. Path enumeration itself can be time consuming, and this approach does not easily generalize to projects that have resources shared between activities, for instance. Therefore, we also present a discrete-event representation which is more complicated, but also more general.

4.4.1 Maximum Path Simulation of the SAN

Figure 4.12 shows a VBASim implementation of the algorithm displayed in Sect. 3.4 and repeated here:

1. set $s = 0$
2. repeat n times:
 - a. generate X_1, X_2, \dots, X_5
 - b. set $Y = \max\{X_1 + X_4, X_1 + X_3 + X_5, X_2 + X_5\}$
 - c. if $Y > t_p$ then set $s = s + 1$
3. estimate θ by $\hat{\theta} = s/n$

Since $\Pr\{Y \leq t_p\}$ is known for this example (see Eq. (3.12)), the true $\theta = \Pr\{Y > t_p\} = 0.165329707$ when $t_p = 5$ is also computed by the program so that we can compare it to the simulation estimate. Of course, in a practical problem we would not know the answer, and we would be wasting our time simulating it if we did. Notice that all of the digits in this probability are correct—assuming that the numerical functions in VBA did their job—although certainly not practically useful.

```

Private Sub SANMax()
' Simple simulation of the SAN using the max function

Dim N As Double, c As Double, tp As Double, _
    Y As Double, Theta As Double
Dim X(1 To 5) As Double
Dim Rep As Integer, i As Integer

Call Report("Pr{Y > tp}", "SANMax", 1, 1)
Call Report("True Theta", "SANMax", 1, 2)
Call InitializeRNSeed
N = 1000
c = 0
tp = 5

For Rep = 1 To N
    For i = 1 To 5
        X(i) = Expon(1, 7)
    Next i
    Y = WorksheetFunction.Max(X(1) + X(4), X(1) + X(3) + X(5), _
        X(2) + X(5))

    If Y > tp Then
        c = c + 1
    End If
Next Rep

Call Report(c / N, "SANMax", 2, 1)

Theta = 1 - ((tp ^ 2 / 2 - 3 * tp - 3) * VBA.Exp(-2 * tp) _
    + (-tp ^ 2 / 2 - 3 * tp + 3) * VBA.Exp(-tp) + 1 - VBA.Exp(-3 * tp))
Call Report(Theta, "SANMax", 2, 2)

End
End Sub

```

Fig. 4.12 Simulation of the SAN as the maximum path through the network

The simulation estimate turns out to be $\hat{\theta} = 0.163$. A nice feature of a probability estimate that is based on i.i.d. outputs is that an estimate of its standard errors is easily computed:

$$\widehat{se} = \sqrt{\frac{\hat{\theta}(1 - \hat{\theta})}{n}}.$$

Thus, \widehat{se} is approximately 0.011, and the simulation has done its job since the true value θ is well within $\pm 1.96\widehat{se}$ of $\hat{\theta}$. This is a reminder that simulations do not deliver *the answer*, such as Eq. (3.12), but do provide the capability to estimate the simulation error, and to reduce that error to an acceptable level by increasing the simulation effort (number of replications).

4.4.2 Discrete-Event Simulation of the SAN

This section uses more advanced VBA constructs than any other part of the book and may be skipped without loss of continuity.

As noted in Sect. 3.4, we can think of the completion of a project activity as an event, and when all of the inbound activities $\mathcal{I}(j)$ to a milestone j are completed then the outbound activities $i \in \mathcal{O}(j)$ can be scheduled, where the destination milestone of activity i is $\mathcal{D}(i)$. Thus, the following generic milestone event is the only one needed:

event milestone (activity ℓ inbound to node j)

$\mathcal{I}(j) = \mathcal{I}(j) - \ell$

if $\mathcal{I}(j) = \emptyset$ then

 for each activity $i \in \mathcal{O}(j)$

 schedule milestone(activity i inbound to node $\mathcal{D}(i)$ at X_i time units later)

end if

Of course, this approach shifts the effort from enumerating all of the paths through the SAN to creating the sets $\mathcal{I}, \mathcal{O}, \mathcal{D}$, but these sets have to be either explicitly or implicitly defined to define the project itself. The key lesson from this example, which applies to many simulations, is that it is possible to program a single event routine to handle many simulation events that are conceptually distinct, and this is done by passing event-specific information to the event routine. In this case we need to pass the inbound activity and the target node, and since this information is needed when the event is executed, not when it is scheduled, we need to store it with the event notice. To do this, we will create a new class module and make use of a feature of the EventNotice class module.

We define a new class module Activity which has two properties: WhichActivity and WhichNode:

```
' Object to model an activity-destination node pair
Public WhichActivity As Integer
Public WhichNode As Integer
```

This object represents an activity, which has a number $1, 2, \dots, 5$, and also a destination node, which we will number with $j = 1$ for a , $j = 2$ for b , $j = 3$ for c and $j = 4$ for d .

When we schedule an activity to be completed at some time in the future, we will associate an Activity with the EventNotice using its WhichObject property; the EventNotice class module is reproduced below:

```
' This is a generic EventNotice object with EventTime,
' EventType and WhichObject attributes
Public EventTime As Double
Public EventType As String
Public WhichObject As Object
```

```
' Add additional problem specific attributes here
```

Perhaps the most complex feature of this simulation is how nodes are represented as a two-dimensional array of VBA collections. A collection is a generic VBA data structure, very much like a one-dimensional array, except that collections may contain objects, not just numbers, and VBA provides methods for adding, removing and referencing elements of a collection. In this case, each element of the array `Nodes(i, j)` is a list of activities, if $i = 1$ it is inbound activities, and if $i = 2$ it is outbound activities, for node j . Thus, `Nodes(1, j)` plays the role of the set $\mathcal{I}(j)$, while `Nodes(2, j)` corresponds to the set $\mathcal{O}(j)$. Another collection called `Destination` associates each activity with its destination and plays the role of the set \mathcal{D} . The key initializations that define the SAN take place in `Sub SANInit`, shown in Fig. 4.13.

One other VBA feature that we use extensively in this example is the ability to define constants, meaning variables whose values are fixed once and for all in the simulation. This is done to make the simulation more readable. For instance,

```
Const a As Integer = 1, b As Integer = 2, _
      c As Integer = 3, d As Integer = 4
Const InTo As Integer = 1, OutOf As Integer = 2
```

By doing this we avoid the need to remember that node c corresponds to the number 3, and what index corresponds to incoming vs. outgoing activities to a node.

Notice (see Fig. 4.14) that the simulation ends when there are no additional activities remaining to be completed. This can be checked via `Calendar.N`, a method that provides the number of events currently on the event calendar. The single milestone event is shown in Fig. 4.15.

A difference between this implementation of the SAN simulation and the one in Sect. 4.4.1 is that here we write out the actual time the project completes on each replication. By doing so, we can estimate $\Pr\{Y > t_p\}$ for any value of t_p by sorting the data and counting how many out of 1,000 replications were greater than t_p . Figure 4.16 shows the empirical cdf of the 1,000 project completion times, which is the simulation estimate of Eq. (3.12).

4.4.3 Issues and Extensions

1. In real projects there are not only activities, but also limited and often shared resources that are needed to complete the activities. Further, there may be specific resource allocation rules when multiple activities contend for the same resource. How might this be modeled in `VBASim`?
2. Time to complete the project is an important overall measure, but at the planning stage it may be more important to discover which activities or resources are the most critical to on-time completion of the project. What additional output measures might be useful for deciding which activities are “critical”?

```
Private Sub SANInit()  
    ' destinations  
    ' Destination(i) corresponds to the destination of activity i  
    Destination.Add b  
    Destination.Add c  
    Destination.Add c  
    Destination.Add d  
    Destination.Add d  
    Dim Inbound As New Collection  
    Dim Outbound As New Collection  
  
    ' node a  
    Outbound.Add 1  
    Outbound.Add 2  
    Set Nodes(InTo, a) = Inbound  
    Set Nodes(OutOf, a) = Outbound  
    Set Outbound = Nothing  
    Set Inbound = Nothing  
  
    ' node b  
    Inbound.Add 1  
    Outbound.Add 3  
    Outbound.Add 4  
    Set Nodes(InTo, b) = Inbound  
    Set Nodes(OutOf, b) = Outbound  
    Set Inbound = Nothing  
    Set Outbound = Nothing  
  
    ' node c  
    Inbound.Add 2  
    Inbound.Add 3  
    Outbound.Add 5  
    Set Nodes(InTo, c) = Inbound  
    Set Nodes(OutOf, c) = Outbound  
    Set Inbound = Nothing  
    Set Outbound = Nothing  
  
    ' node d  
    Inbound.Add 4  
    Inbound.Add 5  
    Set Nodes(InTo, d) = Inbound  
    Set Nodes(OutOf, d) = Outbound  
    Set Inbound = Nothing  
    Set Outbound = Nothing  
  
End Sub
```

Fig. 4.13 Network description for the discrete-event SAN simulation

```

' SAN Simulation using a discrete-event approach
' Nodes is a two-dimensional array of Collections, where each
' Collection is a list of inbound or outbound activities
' to that node
' For Nodes(i, j) = inbound i=1 or outbound i=2
' node j = 1 for a, j = 2 for b, j = 3 for c, j = 4 for d

Dim Nodes(1 To 2, 1 To 4) As Object
Dim Destination As New Collection
Const a As Integer = 1, b As Integer = 2, c As Integer = 3, _
    d As Integer = 4
Const InTo As Integer = 1, OutOf As Integer = 2

Private Sub SAN()
Dim Rep As Integer
Dim NextEvent As EventNotice
Dim ThisActivity As Activity

Call MyInit

For Rep = 1 To 1000
    Call VBASimInit
    Call SANInit ' initializes the activities in the SAN
    Call Milestone(0, a) ' causes outbound activities of node a
                        ' to be scheduled

    Do
        Set NextEvent = Calendar.Remove
        Clock = NextEvent.EventTime
        Set ThisActivity = NextEvent.WhichObject
        Call Milestone(ThisActivity.WhichActivity, _
            ThisActivity.WhichNode)
    Loop Until Calendar.N = 0 ' stop when event calendar is empty

    Call Report(Clock, "SAN", Rep + 1, 1)

Next Rep
End
End Sub

Private Sub MyInit()
Call InitializeRNSeed
Call Report("Completion Time", "SAN", 1, 1)
End Sub

```

Fig. 4.14 Main program for the discrete-event SAN simulation

```

Private Sub Milestone(ActIn As Integer, Node As Integer)

Dim ActOut As Integer, Incoming As Integer, m As Integer
Dim Inbound As Collection
Dim Outbound As Collection

Set Inbound = Nodes(InTo, Node)
Set Outbound = Nodes(OutOf, Node)
m = Inbound.Count

For Incoming = 1 To m
    If Inbound(Incoming) = ActIn Then
        Inbound.Remove (Incoming)
        Exit For
    End If
Next Incoming
Set Nodes(InTo, Node) = Inbound

If Inbound.Count = 0 Then
    m = Outbound.Count
    For ActOut = 1 To m
        Dim ThisActivity As New Activity
        ThisActivity.WhichActivity = Outbound(1)
        ThisActivity.WhichNode = Destination(Outbound(1))
        Call SchedulePlus("Milestone", Expon(1, 1),
            ThisActivity)
        Set ThisActivity = Nothing
        Outbound.Remove (1)
    Next ActOut
End If

End Sub

```

Fig. 4.15 Milestone event for the discrete-event SAN simulation

4.5 Simulating the Asian Option

Here we consider estimating the value of an Asian option

$$v = E \left[e^{-rT} (\bar{X}(T) - K)^+ \right]$$

as described in Sect. 3.5, where the maturity is $T = 1$ year, the risk-free interest rate is $r = 0.05$ and the strike price is $K = \$55$. The underlying asset has an initial value of $X(0) = \$50$ and the volatility is $\sigma^2 = (0.3)^2$. Recall that the key quantity is

$$\bar{X}(T) = \frac{1}{T} \int_0^T X(t) dt$$

the time average of a continuous-time, continuous-state geometric Brownian motion process which we cannot truly simulate on a digital computer. Thus, we approximate

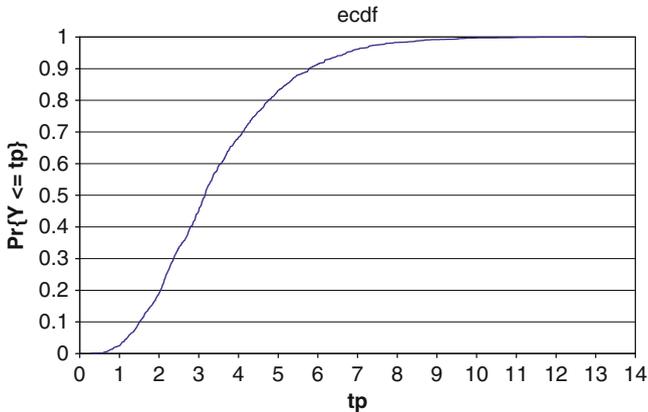


Fig. 4.16 Empirical cdf of the project completion times

it by dividing the interval $[0, T]$ into m steps of size $\Delta t = T/m$ and using the discrete approximation

$$\widehat{X}(T) = \frac{1}{m} \sum_{i=1}^m X(i\Delta t).$$

This makes simulation possible, since

$$X(t_{i+1}) = X(t_i) \exp \left\{ \left(r - \frac{1}{2} \sigma^2 \right) (t_{i+1} - t_i) + \sigma \sqrt{t_{i+1} - t_i} Z_{i+1} \right\}$$

for any increasing sequence of times $\{t_0, t_1, \dots, t_m\}$, where Z_1, Z_2, \dots, Z_m are i.i.d. $N(0, 1)$.

Figure 4.17 is VBA code that uses $m = 32$ steps in the approximation, and makes 10,000 replications to estimate v . Discrete-event structure would slow execution without any obvious benefit, so a simple loop is used to advance time. The value of the option from each replication is written to Excel for post-simulation analysis.

The estimated value of v is \$2.20 with a relative error of just over 2% (recall that the relative error is the standard error divided by the mean). As the histogram in Fig. 4.18 shows, the option is frequently worthless (approximately 68% of the time), but the average payoff, conditional on the payoff being positive, is approximately \$6.95.

4.6 Case Study: Service Center Simulation

This section presents a simulation case based on a project provided by a former student. While still relatively simple, it is more complex than the previous stylized examples, and the answer is not known without simulating. The purpose of this

```

Private Sub AsianOption()
' Simulation of an Asian Option

Dim Replications As Integer
Dim Maturity As Double, Steps As Double
Dim i As Integer, j As Integer
Dim X As Double, Z As Double, InitialValue As Double
Dim Sigma As Double, Sigma2 As Double
Dim Sum As Double, Value As Double
Dim Interval As Double
Dim InterestRate As Double, StrikePrice As Double

' Set option parameters
Replications = 10000
Maturity = 1
Steps = 32
Sigma = 0.3
InterestRate = 0.05
InitialValue = 50
StrikePrice = 55
Interval = Maturity / Steps
Sigma2 = Sigma * Sigma / 2

Call Report("Option Value", "Option", 1, 1)
Call InitializeRNSeed
For i = 1 To Replications
    Sum = 0
    X = InitialValue
    For j = 1 To Steps
        Z = Normal(0, 1, 12)
        X = X * VBA.Exp((InterestRate - Sigma2) * Interval + _
            Sigma * VBA.Sqrt(Interval) * Z)
        Sum = Sum + X
    Next j
    Value = VBA.Exp(-InterestRate * Maturity) * _
        WorksheetFunction.Max(Sum / Steps - StrikePrice, 0)
    Call Report(Value, "Option", i + 1, 1)
Next i

End
End Sub

```

Fig. 4.17 VBA simulation of the Asian option problem

section is to illustrate how one might attack simulation modeling and programming for a realistic problem.

Example 4.1 (Fax center staffing). A service center receives faxed orders throughout the day, with the rate of arrival varying hour by hour. The arrivals are modeled by a nonstationary Poisson process with the rates shown in Table 4.3.

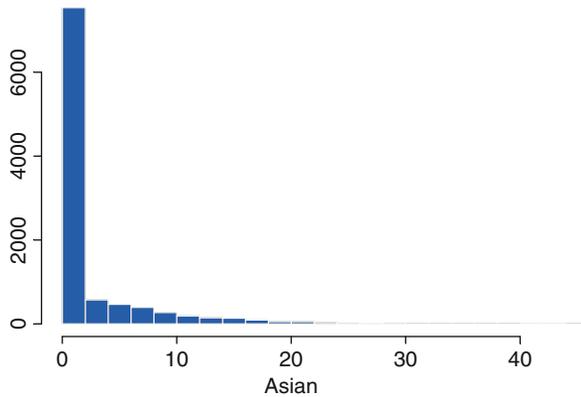


Fig. 4.18 Histogram of the realized value of the Asian option from 10,000 replications

Table 4.3 Arrival rate of faxes by hour

Time	Rate (faxes/minute)
8 a.m.–9 a.m.	4.37
9 a.m.–10 a.m.	6.24
10 a.m.–11 a.m.	5.29
11 a.m.–12 p.m.	2.97
12 p.m.–1 p.m.	2.03
1 p.m.–2 p.m.	2.79
2 p.m.–3 p.m.	2.36
3 p.m.–4 p.m.	1.04

A team of Entry Agents select faxes on a first-come-first-served basis from the fax queue. Their time to process a fax is modeled as normally distributed with mean 2.5 min and standard deviation 1 min. There are two possible outcomes after the Entry Agent finishes processing a fax: either it was a simple fax and the work on it is complete, or it was not simple and it needs to go to a Specialist for further processing. Over the course of a day, approximately 20% of the faxes require a Specialist. The time for a Specialist to process a fax is modeled as normally distributed with mean 4.0 min and standard deviation 1 min.

Minimizing the number of staff minimizes cost, but certain service-level requirements must be achieved. In particular, 96% of all simple faxes should be completed within 10 min of their arrival, while 80% of faxes requiring a Specialist should also be completed (by both the Entry Agent and the Specialist) within 10 min of their arrival.

The service center is open from 8 a.m. to 4 p.m. daily, and it is possible to change the staffing level at 12 p.m. Thus, a staffing policy consists of four numbers: the number of Entry Agents and Specialists before noon, and the number of Entry Agents and Specialists after noon. Any fax that starts its processing before noon completes processing by that same agent before the agent goes off duty; and faxes

in the queues at the end of the day are processed before the agents leave work, and therefore are not carried over to the next day.

The first step in building any simulation model is deciding what question or questions that the model should answer. Knowing the questions helps identify the system performance measures that the simulation needs to estimate, which in turn drives the scope and level of detail in the simulation model.

The grand question for the service center is, what is the minimum number of Entry Agents and Specialists needed for both time periods to meet the service-level requirements? Therefore, the simulation must at least provide an estimate of the percentage of faxes of each type entered within 10 min, given a specific staff assignment.

Even when there seems to be a clear overall objective (e.g., minimize the staff required to achieve the service-level requirement), we often want to consider trade-offs around that objective. For instance, if meeting the requirement requires a staff that is so large that they are frequently underutilized, or if employing the minimal staff means that the Entry Agents or Specialists frequently have to work well past the end of the day, then we might be willing to alter the service requirement a bit. Statistics on the number and the time spent by faxes in queue, and when the last fax of each day is actually completed, provide this information. Including additional measures of system performance, beyond the most critical ones, makes the simulation more useful.

Many discrete-event, stochastic simulations involve entities that dynamically flow through some sort of queueing network where they compete for resources. In such simulations, identifying the entities and resources is a good place to start the model. For this service center the faxes are clearly the dynamic entities, while the Entry Agents and Specialists are resources. The fax machines themselves might also be considered a resource, especially if they are heavily utilized or if outgoing as well as incoming faxes use the same machines. It turns out that for this service center there is a bank of fax machines dedicated to incoming faxes, so it is reasonable to treat the arrival of faxes as an unconstrained external arrival process. This fact was not stated in the original description of the problem; follow-up questions are often needed to fully understand the system of interest.

Whenever there are scarce resources, queues may form. Queues are often first-in-first-out, with one queue for each resource, as they are in this service center. However, queues may have priorities, and multiple queues may be served by the same resource, or a single queue may feed multiple resources. Queueing behavior is often a critical part of the model.

When the simulation involves entities flowing through a network of queues, then there can be two types of arrivals: arrivals from outside of the network and arrivals internal to the network. Outside arrivals are like those we have already seen in the $M(t)/M/\infty$ and $M/G/1$ examples. Internal arrivals are departures from one queue that become arrivals to others. How these are modeled depends largely on whether the departure from one queue is an immediate arrival to the next—in which case the departure and arrival events are effectively the same thing—or whether there is some sort of transportation delay—in which case the arrival to the next queue should be

scheduled as a distinct event. For the service center the arrival of faxes to the Entry Agents is an outside arrival process, while the 20% of faxes that require a Specialist are internal arrivals from the Entry Agents to the Specialists.

Critical to experiment design is defining what constitutes a replication. Replications should be independent and identically distributed. Since the service center does not carry faxes over from 1 day to the next, a “day” defines a replication. If faxes did carry over, but all faxes are cleared weekly, then a replication might be defined by a work week. However, if there is always significant carry over from 1 day to the next, then a replication might have to be defined arbitrarily.

The work day at the service center is 8 h; however the staff does not leave until all faxes that arrive before 4 p.m. are processed. If we defined a replication to be exactly 8 h then we could be fooled by a staffing policy that allows a large queue of faxes to build up toward the end of the day, since the entry of those faxes would not be included in our statistics. To model a replication that ends when there is no additional work remaining, we will cut off the fax arrivals at 4 p.m. and then end the simulation when the event calendar is empty. This works because idle Entry Agents and Specialists will always take a fax from their queue if one is available.

Rather than walk through the VBASim code line by line, we will point out some highlights to facilitate the reader’s understanding of the code.

Figure 4.19 shows the global declarations for the service center simulation. Of particular note are the two `DTStat` statements defining `Regular10` and `Special10`. These will be used to obtain the fraction of regular and special faxes that are processed within the 10-min requirement by recording a 1 for any fax that meets the requirement, and a 0 otherwise. The mean of these values is the desired fraction.

The main program for the simulation is in Fig. 4.20. Of particular importance is the condition that ends the main simulation loop:

```
Loop Until Calendar.N = 0
```

The `N` method of `Calendar` returns the current number of pending events. When the event calendar is empty, then there are no additional faxes to process, and no pending arrival of a fax. This condition will only hold after 4 p.m. and once all remaining faxes have been entered.

Also notice the event `ChangeStaff`, which is scheduled to occur at noon (240 min). Here we use the `SetUnits` method of the `Resource` to change the staffing levels. The reader should look at the VBASim `Resource` class object and convince themselves that even if we reduce the number of staff at noon the faxes in process will not be affected.

The staffing policy to be simulated is read in from four cells of the worksheet `Fax`. By placing a VBA control button to start the simulation on the worksheet we can simulate different staffing policies, and view the results, without actually touching the VBA code.

The function `NSPP_Fax` generates the interarrival times for faxes with the desired time-varying rate; this function will be described in Chap. 6.

```

' FaxCenter Simulation
' Parameters we may want to change
Dim MeanRegular As Double      ' mean entry time regular faxes
Dim VarRegular As Double      ' variance entry time regular faxes
Dim MeanSpecial As Double     ' mean entry time special faxes
Dim VarSpecial As Double     ' variance entry time special faxes
Dim RunLength As Double       ' length of the working day
Dim NumAgents As Integer      ' number of regular agents
Dim NumSpecialists As Integer ' number of special agents
Dim NumAgentsPM As Integer    ' number of regular agents after noon
Dim NumSpecialistsPM As Integer ' number of special agents after noon

' Global objects needed for simulation
Dim RegularQ As New FIFOQueue ' queue for all faxes
Dim SpecialQ As New FIFOQueue ' queue for special faxes
Dim RegularWait As New DTStat ' discrete-time statistics on
                               ' fax waiting
Dim SpecialWait As New DTStat ' discrete-time statistics on
                               ' special fax waiting
Dim Regular10 As New DTStat   ' discrete-time statistics on
                               ' < 10 minutes threshold
Dim Special10 As New DTStat   ' discrete-time statistics on
                               ' < 10 minutes threshold
Dim Agents As New Resource    ' entry agents resource
Dim Specialists As New Resource ' specialists resource
Dim ARate(1 To 8) As Double   ' arrival rates
Dim MaxRate As Double         ' maximum arrival rate
Dim Period As Double          ' period for which arrival rate
                               ' stays constant
Dim NPeriods As Integer       ' number of periods in a "day"

```

Fig. 4.19 Global declarations for service center simulation

Figure 4.21 contains the arrival and end-of-entry events for faxes at the Entry Agents. The next arrival is scheduled only if `Clock < RunLength`; in this way we cut off fax arrivals after 4p.m. The `EndOfEntry` event passes 20% of the faxes directly and immediately to the Specialists by executing `Call SpecialArrival(DepartingFax)` and passing the pointer to the `DepartingFax` entity. Equivalently, we could have scheduled a `SpecialArrival` event to occur zero time units into the future (or nonzero time units if it takes time to transport the fax).

The `Record` method of the `DTStat Regular10` is used to collect a 0 or 1 depending on whether the total wait was less than 10 min.

The arrival and end-of-entry events for the Specialists, shown in Fig. 4.22, work similarly to those of the Entry Agents.

Initializations that occur once are shown in Fig. 4.23.

Ten replications of this simulation with a staffing policy of 15 Entry Agents in the morning and 9 in the afternoon, and 6 Specialists in the morning and 3 in the afternoon, gives 0.98 ± 0.02 for the fraction of regular faxes entered in 10 min or

```

Public Sub FaxCenterSim()
    Dim Reps As Integer
    Dim NextEvent As EventNotice
    ' Read in staffing policy
    NumAgents = Worksheets("Fax").Cells(25, 5)
    NumAgentsPM = Worksheets("Fax").Cells(25, 6)
    NumSpecialists = Worksheets("Fax").Cells(26, 5)
    NumSpecialistsPM = Worksheets("Fax").Cells(26, 6)
    Call MyInit
    For Reps = 1 To 10
        Call VBASimInit
        Agents.SetUnits (NumAgents)
        Specialists.SetUnits (NumSpecialists)
        Call Schedule("Arrival", NSPP_Fax(ARate, MaxRate, NPeriods, _
            Period, 1))
        Call Schedule("ChangeStaff", 4 * 60)
        Do
            Set NextEvent = Calendar.Remove
            Clock = NextEvent.EventTime
            Select Case NextEvent.EventType
                Case "Arrival"
                    Call Arrival
                Case "EndOfEntry"
                    Call EndOfEntry(NextEvent.WhichObject)
                Case "EndOfEntrySpecial"
                    Call EndOfEntrySpecial(NextEvent.WhichObject)
                Case "ChangeStaff"
                    Agents.SetUnits (NumAgentsPM)
                    Specialists.SetUnits (NumSpecialistsPM)
            End Select
        Loop Until Calendar.N = 0 ' stop when event calendar empty
    ' Write output report for each replication
    Call Report(RegularWait.Mean, "Fax", Reps + 1, 1)
    Call Report(RegularQ.Mean, "Fax", Reps + 1, 2)
    Call Report(Agents.Mean, "Fax", Reps + 1, 3)
    Call Report(SpecialWait.Mean, "Fax", Reps + 1, 4)
    Call Report(SpecialQ.Mean, "Fax", Reps + 1, 5)
    Call Report(Specialists.Mean, "Fax", Reps + 1, 6)
    Call Report(Regular10.Mean, "Fax", Reps + 1, 7)
    Call Report(Special10.Mean, "Fax", Reps + 1, 8)
    Call Report(Clock, "Fax", Reps + 1, 9)
    Next Reps
End Sub

```

Fig. 4.20 Main program for service center simulation

```

Private Sub Arrival()
' Schedule next fax arrival if < 4 PM
  If Clock < RunLength Then
    Call Schedule("Arrival", NSPP_Fax(ARate, MaxRate, _
      NPeriods, Period, 1))
  Else
    Exit Sub
  End If
' Process the newly arriving Fax
  Dim Fax As New Entity
  If Agents.Busy < Agents.NumberOfUnits Then
    Agents.Seize (1)
    Call SchedulePlus("EndOfEntry", Normal(MeanRegular, _
      VarRegular, 2), Fax)
  Else
    RegularQ.Add Fax
  End If
  Set Fax = Nothing
End Sub

Private Sub EndOfEntry(DepartingFax As Entity)
  Dim Wait As Double
' Record wait time if regular; move on if special
  If Uniform(0, 1, 3) < 0.2 Then
    Call SpecialArrival(DepartingFax)
  Else
    Wait = Clock - DepartingFax.CreateTime
    RegularWait.Record (Wait)
    If Wait < 10 Then
      Regular10.Record (1)
    Else
      Regular10.Record (0)
    End If
  End If
  Set DepartingFax = Nothing
' Check to see if there is another Fax; if yes start entry
' otherwise free the agent
  If RegularQ.NumQueue > 0 And _
    Agents.NumberOfUnits >= Agents.Busy Then
    Set DepartingFax = RegularQ.Remove
    Call SchedulePlus("EndOfEntry", Normal(MeanRegular, _
      VarRegular, 2), DepartingFax)
    Set DepartingFax = Nothing
  Else
    Agents.Free (1)
  End If
End Sub

```

Fig. 4.21 Events for Entry Agents

```

Private Sub SpecialArrival(SpecialFax As Entity)
' If special agent available, start entry by seizing
' the special agent
  If Specialists.Busy < Specialists.NumberOfUnits Then
    Specialists.Seize (1)
    Call SchedulePlus("EndOfEntrySpecial", Normal(MeanSpecial, _
      VarSpecial, 4), SpecialFax)
  Else
    SpecialQ.Add SpecialFax
  End If
  Set SpecialFax = Nothing
End Sub

Private Sub EndOfEntrySpecial(DepartingFax As Entity)
  Dim Wait As Double
' Record wait time and indicator if < 10 minutes
  Wait = Clock - DepartingFax.CreateTime
  SpecialWait.Record (Wait)
  If Wait < 10 Then
    Special10.Record (1)
  Else
    Special10.Record (0)
  End If
  Set DepartingFax = Nothing
' Check to see if there is another Fax; if yes start entry
' otherwise free the specialist
  If SpecialQ.NumQueue > 0 And _
    Specialists.NumberOfUnits >= Specialists.Busy Then
    Set DepartingFax = SpecialQ.Remove
    Call SchedulePlus("EndOfEntrySpecial", Normal(MeanSpecial, _
      VarSpecial, 4), DepartingFax)
    Set DepartingFax = Nothing
  Else
    Specialists.Free (1)
  End If
End Sub

```

Fig. 4.22 Events for Specialists

less, and 0.81 ± 0.06 for the special faxes. The “ \pm ” are 95% confidence intervals. This policy appears to be close to the requirements, although if we absolutely insist on 80% for the special faxes then additional replications are needed to narrow the confidence interval.

4.6.1 Issues and Extensions

1. There are many similarities between the programming for this simulation and the event-based simulation of the $M/G/1$ queue. However, there is an important

```

Private Sub MyInit ()
    Call InitializeRNSeed
    MeanRegular = 2.5
    VarRegular = 1#
    MeanSpecial = 4
    VarSpecial = 1#
    RunLength = 480
' Add queues, resources and statistics that need to be
' initialized between replications to the global collections
    TheDTStats.Add RegularWait
    TheDTStats.Add SpecialWait
    TheDTStats.Add Regular10
    TheDTStats.Add Special10
    TheQueues.Add RegularQ
    TheQueues.Add SpecialQ
    TheResources.Add Agents
    TheResources.Add Specialists
' Write headings for the output reports
    Call Report("Ave Reg Wait", "Fax", 1, 1)
    Call Report("Ave Num Reg Q", "Fax", 1, 2)
    Call Report("Agents Busy", "Fax", 1, 3)
    Call Report("Ave Spec Wait", "Fax", 1, 4)
    Call Report("Ave Num Spec Q", "Fax", 1, 5)
    Call Report("Specialists Busy", "Fax", 1, 6)
    Call Report("Reg < 10", "Fax", 1, 7)
    Call Report("Spec < 10", "Fax", 1, 8)
    Call Report("End Time", "Fax", 1, 9)
' Arrival process data
    NPeriods = 8
    Period = 60
    MaxRate = 6.24
    ARate(1) = 4.37
    ARate(2) = 6.24
    ARate(3) = 5.29
    ARate(4) = 2.97
    ARate(5) = 2.03
    ARate(6) = 2.79
    ARate(7) = 2.36
    ARate(8) = 1.04
End Sub

```

Fig. 4.23 Initializations for service center simulation

difference that is due to having multiple agents. For the $M/G/1$ queue, a single `FIFOQueue` object holds both the customer in service (who is at the front of the queue) and the customers waiting for service. This approach does not work for the Fax Center because when there are multiple agents the faxes need not complete entry in the same order in which they arrive. To accommodate this, the `FIFOQueue` holds only those faxes waiting for entry, and the `Entity` representing a fax that is being entered is stored with the `EventNotice` for the end-of-entry event. This is accomplished by the statement

```
Call SchedulePlus("EndOfEntry", _
    Normal(MeanRegular, VarRegular, 2), DepartingFax)
```

SchedulePlus allows an object (DepartingFax in this case) to be assigned to the .WhichObject property of the EventNotice. The Entity can then be passed to the event using the statement

```
Call EndOfEntry(NextEvent.WhichObject)
```

2. The fax entry times were modeled as being normally distributed. However, the normal distribution admits negative values, which certainly does not make sense. What should be done about this? Consider mapping negative values to 0, or generating a new value whenever a negative value occurs. Which is more likely to be realistic and why?

Exercises

1. Simulate an $M(t)/G/\infty$ queue where G corresponds to an Erlang distribution with fixed mean but try different numbers of phases. That is, keep the mean service time fixed but change the variability. Is the expected number in queue sensitive to the variance in the service time?
2. This problem assumes a more advanced background in stochastic processes. In the simulation of the $M(t)/M/\infty$ queue there could be a very large number of events on the event calendar: one "Arrival" and one "Departure" for *each* car currently in the garage. However, properties of the exponential distribution can reduce this to no more than two events. Let $\beta = 1/\tau$ be the departure rate for a car (recall that τ is the mean parking time). If at any time we observe that there are N cars in the garage (no matter how long they have been there), then the time until the first of these cars departs is exponentially distributed with mean $1/(N\beta)$. Use this insight to build an $M(t)/M/\infty$ simulation with at most two pending events, next arrival and next departure. *Hint*: Whenever an arrival occurs the distribution of the time until the next departure changes, so the scheduled next departure time must again be generated.
3. For the hospital problem, simulate the current system in which the receptionist's service time is well modeled as having an Erlang-4 distribution with mean 0.6 min. Compare the waiting time to the proposed electronic kiosk alternative.
4. Beginning with the event-based $M/G/1$ simulation, implement the changes necessary to make it an $M/G/s$ simulation (a single queue with any number of servers). Keeping $\lambda = 1$ and $\tau/s = 0.8$, simulate $s = 1, 2, 3$ servers and compare the results. What you are doing is comparing queues with the same service capacity, but with 1 fast server as compared to two or more slower servers. State clearly what you observe.
5. Modify the VBASim event-based simulation of the $M/G/1$ queue to simulate an $M/G/1/c$ retrial queue. This means that customers who arrive to find c customers in the system (including the customer in service) leave immediately,

but arrive again after an exponentially distributed amount of time with mean MeanTR. *Hint:* The existence of retrial customers should not affect the arrival process for first-time arrivals.

6. Modify the SAN simulation to allow each activity to have a different mean time to complete (currently they all have mean time 1). Use a collection to hold these mean times.
7. Try the following numbers of steps for approximating the value of the Asian option to see how sensitive the value is to the step size: $m = 8, 16, 32, 64, 128$.
8. In the simulation of the Asian option, the sample mean of 10,000 replications was 2.198270479, and the standard deviation was 4.770393202. Approximately how many replications would it take to decrease the relative error to less than 1%?
9. For the service center, increase the number of replications until you can be confident that suggested policy does or does not achieve the 80% entry in less than 10 min requirement for special faxes.
10. For the service center, find the minimum staffing policy (in terms of total number of staff) that achieves the service-level requirement. Examine the other statistics generated by the simulation to make sure you are satisfied with this policy.
11. For the service center, suppose that Specialists earn twice as much as Entry Agents. Find the minimum cost staffing policy that achieves the service-level requirement. Examine the other statistics generated by the simulation to make sure you are satisfied with this policy.
12. For the service center, suppose that the staffing level can change hourly, but once an Agent or Specialist comes on duty they must work for 4 h. Find the minimum staffing policy (in terms of total number of staff) that achieves the service-level requirement.
13. For the service center, pick a staffing policy that fails to achieve the service level requirements by 20% or more. Rerun the simulation with a replication being defined as exactly 8 h, but do not carry waiting faxes over to the next day. How much do the statistics differ using the two different ways to end a replication?
14. The function NSPP_Fax is listed below. This function implements the thinning method described in Sect. 4.2 for a nonstationary Poisson process with piecewise-constant rate function. Study it and describe how it works.

```
Private Function NSPP_Fax(ARate() As Double, MaxRate As Double, _
    NPeriods As Integer, Period As Double, _
    Stream As Integer) As Double
' This function generates interarrival times from a
' NSPP with piecewise constant arrival rate over a
' fixed time of Period*NPeriod time units

' ARate = array of arrival rates over a common length Period
' MaxRate = maximum value of ARate
' Period = time units between (possible) changes in arrival rate
' NPeriods = number of time periods in ARate
```

```

Dim i As Integer
Dim PossibleArrival As Double
PossibleArrival = Clock + Expon(1 / MaxRate, Stream)
i = WorksheetFunction.Min(NPeriods, _
    WorksheetFunction.Ceiling(PossibleArrival / Period, 1))
Do Until Uniform(0, 1, Stream) < ARate(i) / MaxRate
    PossibleArrival = PossibleArrival + Expon(1 / MaxRate, Stream)
    i = WorksheetFunction.Min(NPeriods, _
        WorksheetFunction.Ceiling(PossibleArrival / Period, 1))
Loop
NSPP_Fax = PossibleArrival - Clock
End Function

```

15. The phone desk for a small office is staffed from 8 a.m. to 4 p.m. by a single operator. Calls arrive according to a Poisson process with rate 6 per hour, and the time to serve a call is uniformly distributed between 5 and 12 min. Callers who find the operator busy are placed on hold, if there is space available, otherwise they receive a busy signal and the call is considered “lost.” In addition, 10% of callers who do not immediately get the operator decide to hang up rather than go on hold; they are not considered lost, since it was their choice. Because the hold queue occupies resources, the company would like to know the smallest capacity (number of callers) for the hold queue that keeps the daily fraction of lost calls under 5%. In addition, they would like to know the long-run utilization of the operator to make sure he or she will not be too busy. Use VBASim to simulate this system and find the required capacity for the hold queue. Model the callers as class *Entity*, the hold queue as class *FIFOQueue* and the operator as class *Resource*. Use the VBASim functions *Expon* and *Uniform* for random-variate generation. Use class *DTStat* to estimate the fraction of calls lost (record a 0 for calls not lost, a 1 for those that are lost so that the sample mean is the fraction lost). Use the statistics collected by class *Resource* to estimate the utilization.
16. Software Made Personal (SMP) customizes software products in two areas: financial tracking and contact management. They currently have a customer support call center that handles technical questions for owners of their software from the hours of 8 a.m. to 4 p.m. Eastern time.

When a customer calls, they first listen to a recording that asks them to select among the product lines; historically 59% are financial products and 41% contact management products. The number of customers who can be connected (talking to an agent or on hold) at any one time is essentially unlimited. Each product line has its own agents. If an appropriate agent is available then the call is immediately routed to the agent; if an appropriate agent is not available, then the caller is placed in a hold queue (and listens to a combination of music and ads). SMP has observed that hang-ups very rarely happen.

SMP is hoping to reduce the total number of agents they need by cross-training agents so that they can answer calls for any product line. Since the agents will not be experts across all products, this is expected to increase the time to process a call by about 5%. The question that SMP has asked you to

answer is how many cross-trained agents are needed to provide service at the same level as the current system.

Incoming calls can be modeled as a Poisson arrival process with a rate of 60 per hour. The mean time required for an agent to answer a question is 5 min, with the actual time being Erlang-2 for financial calls, and Erlang-3 for contact management calls. The current assignment of agents is four for financial and three for contact management. Simulate the system to find out how many agents are needed to deliver the same level of service in the cross-trained system as in the current system.

17. Beginning with the `FIFOQueue` class module, develop a new class module called `PriorityQueue`. When an `Entity` is added to the `PriorityQueue` its position depends on the entity's `Priority` property with high value of `Priority` first. Be sure to add the `Priority` property to the `Entity` class module.

Chapter 5

Two Views of Simulation

Prior to this chapter, the focus of the book has been on building—literally writing computer code for—simulation models. This chapter sets up the remainder of the book which addresses issues related to experiment design and analysis. To do so, we present two different, but complementary ways of viewing computer simulation. Section 5.1 makes clear the roles of real systems, simulated systems, and the conceptual systems that the analyst wants to design. This framework highlights sources of error in a simulation study; it is abstract, but not mathematically formal. Section 5.2 views simulation output as a stochastic process to provide a framework for designing simulation experiments and analyzing the results; it is a mathematical treatment that provides the foundation for statistical analysis. Both of these perspectives are necessary to have a deep understanding of stochastic simulation.

5.1 A Framework for Simulation Modeling and Analysis

Example 5.1 (The call center). A software company has a customer call center for inquiries or problems related to its products. Currently the agents are trained only for particular products, and therefore they only handle calls related to their product. The company would like to reduce staff, while still delivering the same level of service, by cross-training some agents to handle calls regarding more than one product. What staff level is required?

This situation has characteristics that are common to many system-design problems for which we seek a simulation solution: There is a conceptual design for a new system (the call center with cross-trained agents). There is also an existing real system that is related to the conceptual one, but is not identical to it (the current call center with dedicated agents). Using the real system as a starting point, a simulated system will be constructed to evaluate how well the conceptual system might work. In this section we describe a framework for simulation problems such as the call center, a framework that will be informative as we think about simulation input modeling, output analysis and experiment design (the subjects of Chaps. 6–8).

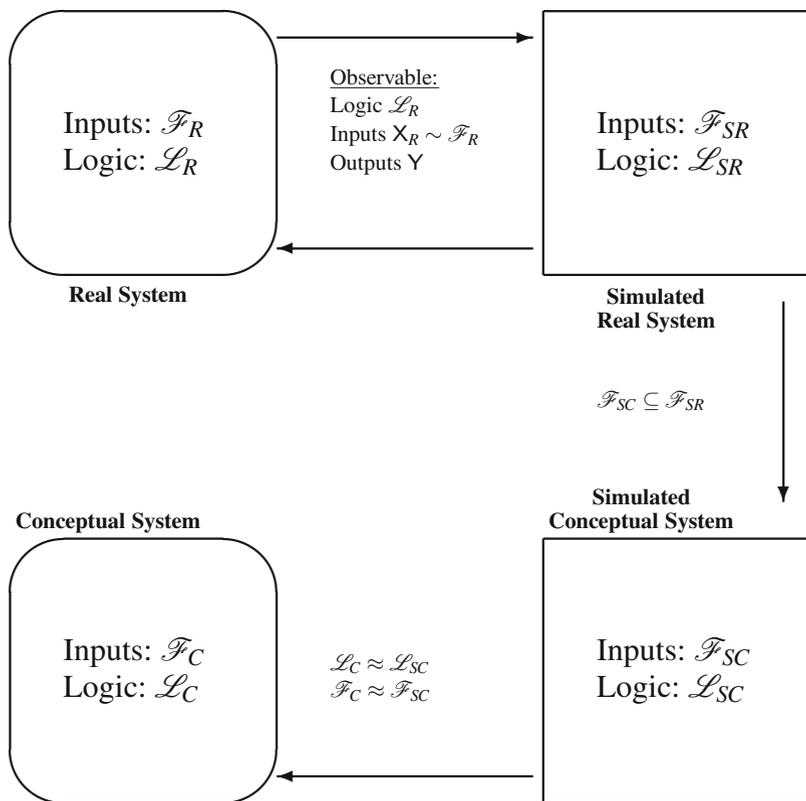


Fig. 5.1 Relationship between real, simulated and conceptual systems.

For the purpose of this framework, a *system* consists of *inputs* and *logic*. Inputs are the uncertain (stochastic) components of a system, while the logic may be thought of as a collection of rules or algorithms that govern how the system behaves as a function of the inputs. Let \mathcal{L} denote logic, and let \mathcal{F} denote the probability model that governs the inputs (often a collection of probability distributions). In our framework, there are three different kinds of systems: Real systems (R), simulated systems (S), and conceptual systems (C); by a “conceptual” system we usually mean a system that could exist, but does not yet exist. Systems have performance parameters that are implied by the inputs and logic; we will denote these generically by μ . Refer to Fig. 5.1 for the discussion that follows.

An *experiment* on a system, denoted by \mathcal{E} , is a recipe for exercising the system to observe its behavior. Experiments on real systems will often be field observations of the system as it currently operates, because a designed experiment on a real system might be expensive, disruptive or dangerous. Simulation experiments, on the other hand, should be designed using statistical principles because the simulated system is completely under our control. Experiments on a conceptual system are, by definition, impossible, which is why we do simulations.

An experiment on a system generates observed behavior, which for our purposes is always numerical. Denote numerical observations of a system's inputs by X and its outputs by Y , both collections of random variables in general. Outputs are derived quantities that are a function of the inputs, logic and experiment design. It is very important to understand that Y is not well defined without an experiment, \mathcal{E} . Observing a real system for a month is different than observing it for a day; running a simulation for ten replications is different than running it for 1,000. This is critical, because performance parameters μ are estimated by functions of the outputs, $\mathcal{T}(Y)$, and therefore properties of these estimators depend on the experiment design.

For example, in the current call center the inputs \mathcal{F}_R include the arrival process of calls, and the time required to resolve a caller's problem (excluding the time they spend on hold before talking to an agent). These are inputs because they are uncertain, stochastic quantities that we do not view as being functions of some more basic stochastic quantities. The call center logic \mathcal{L}_R includes the hours of operation and the rules for how calls are routed to agents. A performance parameter μ_R that measures customer satisfaction is the long-run average time a caller spends on hold before talking to an agent. Therefore, an output of interest Y_R is how long individual callers spend on hold. And an experiment \mathcal{E}_R might be logging data on call center activity (call arrival times, times to resolve problems, and caller hold times) for several months.

In our typical situation, there is a conceptual system $\{\mathcal{L}_C, \mathcal{F}_C\}$ whose performance μ_C we would like to predict; in the example it is the call center with cross-trained agents that we might want to implement. If we could perform an experiment on this system, then we could use observed outputs Y_C to estimate μ_C (e.g., the mean hold time when using cross-trained agents). Of course, we cannot do this directly, because the system is conceptual. We can, however, perform an experiment on the real system $\{\mathcal{L}_R, \mathcal{F}_R\}$; specifically, we can study its logic \mathcal{L}_R , and observe its inputs X_R and outputs Y_R . We do this to aid in developing a simulation of the conceptual system $\{\mathcal{L}_{SC}, \mathcal{F}_{SC}\}$. Notice that while the rules that govern a system's behavior \mathcal{L}_R may be at least partially observable, the input probability model \mathcal{F}_R is not; only the input data itself X_R can be observed. For instance, we can log the times that calls arrive, but not the probability mechanism that generated the calls.

Simulation is useful for evaluating a conceptual system design when we believe we understand the critical aspects of the conceptual system's logic \mathcal{L}_C (often that logic is our choice, like using cross-trained agents), and we also believe that $\mathcal{F}_C \subseteq \mathcal{F}_R$ —that is, the input processes for the conceptual system are a subset of the input processes for the real system. For instance, using cross-trained agents likely will have no impact on the arrival process of calls. This leads naturally to constructing two simulated systems:

- $SR = \{\mathcal{L}_{SR}, \mathcal{F}_{SR}\}$: This is the simulation of the real system (SR stands for “simulated-real” system). Ideally this simulated system has performance parameters μ_{SR} that are close to those of the real system μ_R so we can validate it.
- $SC = \{\mathcal{L}_{SC}, \mathcal{F}_{SC}\}$: This is the simulation of the conceptual system (SC stands for “simulated-conceptual” system). Ideally this simulated system has performance parameters μ_{SC} that are close to those of the conceptual system μ_C . Recall that

the input models for the simulated-conceptual system are a subset of the input models for the simulated-real system, while the conceptual system's logic is our choice.

With this framework, we can now define the simulation design and analysis problems that will occupy the remainder of the book:

Input modeling: This is the process of choosing the simulation input probability models \mathcal{F}_{SR} that approximate the input models of the real system \mathcal{F}_R . When real system inputs X_R can be observed, then we may fit distributions to the data; otherwise we have to use subjective methods. Similarly, input models for the simulated conceptual system \mathcal{F}_{SC} must be chosen.

Variate generation: Given the simulation input models \mathcal{F}_{SR} and \mathcal{F}_{SC} , we need methods to generate realizations from them to drive the simulation. We touched on this topic in Sect. 2.2.1, but will go into more depth and describe variate-generation algorithms for more complex input models in the following chapter.

Validation: In theory, this is the process of deciding whether the performance parameters of the simulated-conceptual system μ_{SC} are close enough to the performance parameters of the conceptual system μ_C for the simulation to be useful for the decisions that it needs to support. Unfortunately, this sort of validation is impossible to achieve by any quantitative means. For this reason we often construct a simulation of the real system, $\{\mathcal{L}_{SR}, \mathcal{F}_{SR}\}$, because we can validate it with respect to the real system $\{\mathcal{L}_R, \mathcal{F}_R\}$. Knowing we can model the real system gives us confidence that our simulation of the conceptual system will also be useful (although this is certainly not a guarantee).

Our definition of validation as a quantitative evaluation of $\{\mathcal{L}_{SR}, \mathcal{F}_{SR}\}$ with respect to $\{\mathcal{L}_R, \mathcal{F}_R\}$ is narrower than many researchers who also provide non-quantitative ways to validate $\{\mathcal{L}_{SC}, \mathcal{F}_{SC}\}$ with respect to $\{\mathcal{L}_C, \mathcal{F}_C\}$. For a broader perspective, see Sargent (2011).

Output analysis: The goal of the simulation experiment is to estimate performance parameters μ_{SC} (which we hope are close to those of the conceptual system μ_C). Output analysis is concerned with designing a simulation experiment \mathcal{E}_{SC} for $\{\mathcal{L}_{SC}, \mathcal{F}_{SC}\}$, and estimating the performance characteristics μ_{SC} via the generated outputs Y_{SC} . If, as is often the case, there is more than one conceptual system, then output analysis may include designing an experiment to identify the best conceptual system.

With this framework in place, the sources of error in solving a problem via simulation can be defined.

Modeling error: $\mathcal{L}_{SC} \neq \mathcal{L}_C$ or $\mathcal{L}_{SR} \neq \mathcal{L}_R$. There is almost always modeling error, since the real or conceptual system will typically be richer and more complicated than the simulation model, with behaviors that just cannot be quantified by a collection of algorithmic rules. For instance, in the call center example the rule for assigning incoming calls to agents may be well defined, while the rule for passing a call along to a supervisor is more ambiguous and agent dependent. The fiction that a real system has quantifiable logic is necessary for modeling, but is

not strictly true. The key is whether the logic we use in our simulation model is accurate enough to obtain useful results. Simulation logic was the focus of Chaps. 1–4.

Input uncertainty: $\mathcal{F}_{SC} \neq \mathcal{F}_C$ or $\mathcal{F}_{SR} \neq \mathcal{F}_R$. The true input distributions \mathcal{F}_R are only discernible through the observed input data X_R , which is always a finite sample. We typically fit \mathcal{F}_{SR} to these data, so there will be error. The problem is compounded when we move to \mathcal{F}_{SC} . Again, the existence of “true” distributions \mathcal{F}_C and \mathcal{F}_R is not strictly correct, but is the basis of all of statistics, and it allows input modeling to be treated as a statistical problem.

Estimation error: $\mathcal{I}(Y_{SC}) \neq \mu_{SC}$. Experiments on simulated systems are statistical experiments because they generate realizations of the output random variables Y_{SC} to estimate performance properties μ_{SC} ; but again, only with a finite sample. Therefore, we need to account for estimation error, often via confidence intervals for the unknown performance parameters. Estimation error is the easiest to quantify, but there can still be complications due to bias and correlated data.

Clearly, not all simulation studies have all the features of the framework described here. For instance, input modeling without any real data is a common occurrence. Also, there may be no real system, only one or more conceptual ones that we want to evaluate. These are special cases of our general framework.

5.2 Simulation as a Stochastic Process

In this section we focus on simulation output processes, the data generated by a simulation that are used to estimate system performance. We think of simulation output as a *stochastic process*: an indexed sequence of random variables defined on a common probability space. Depending on the output process, the index may be time, observation count or replication number. We are particularly interested in large-sample behavior of summary statistics as the index (time, observation count or number of replications) increases. Large-sample, or asymptotic, behavior is relevant for simulation as we are frequently able to simulate very large samples, unlike, for instance, expensive physical experiments. To set up the formal description of asymptotic behavior, we start with a less formal simulation description.

5.2.1 Simulated Asymptotic Behavior

To illustrate the key aspects of asymptotic behavior we will use the $M/G/1$ queue example of Sects. 3.2 and 4.3; recall that this is a single-server queue that was used to model a hospital reception system to estimate long-run average delay in queue, which we will denote by μ . In the simulation we observed Y_1, Y_2, \dots, Y_m , the delay in queue of the first m patients and visitors, and estimated the long-run average by

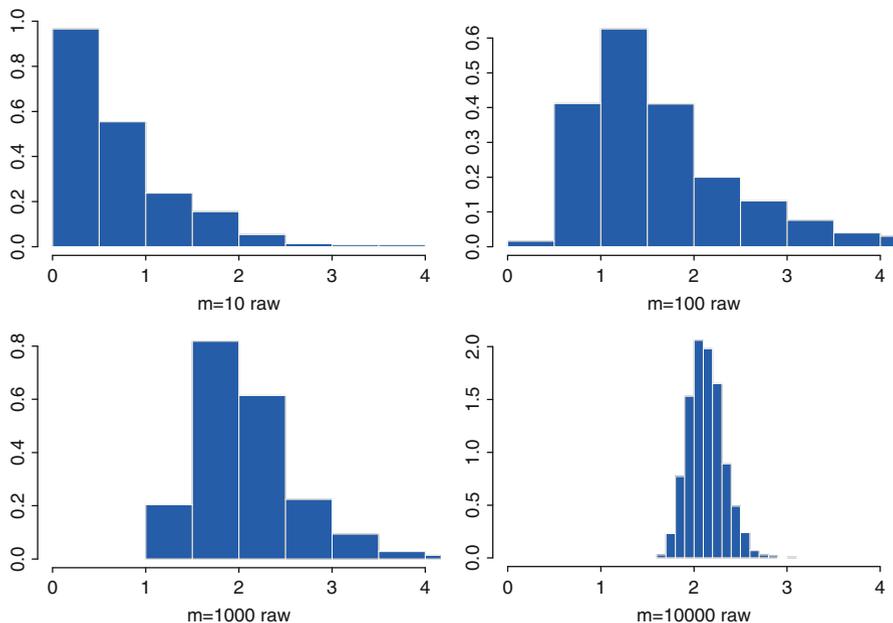


Fig. 5.2 Histogram of the 1,000 replication averages at $m = 10, 100, 1,000$ and $10,000$

the sample mean $\bar{Y}(m) = \sum_{i=1}^m Y_i/m$. We will consider the same case as in Sect. 4.3 where the arrival rate is one customer per minute, and the service time has an Erlang distribution with mean $\tau = 0.8$ min and three phases. For this model the Pollaczek–Khinchine formula in Eq. (3.4) gives $\mu = 2.133$ min, so of course simulation is not actually needed.

Now suppose that you and 999 of your friends are hired by “the boss” to each simulate one replication of this system for $m = 10,000$ patients and visitors, choosing your initial random number seeds independently, randomly and without collaboration. However, because your boss is eager to know the outcome, you will each report your results to her at check points $m = 10, 100, 1,000$ and finally $10,000$ simulated patients and visitors. That is, each of you will report $\bar{Y}(10)$, $\bar{Y}(100)$, $\bar{Y}(1,000)$ and $\bar{Y}(10,000)$. *You* only get to see the result of your one replication, but your boss gets to see them all. We will look at the results both from your and your boss’s perspectives.

Figure 5.2 shows the histograms that your boss would see of the 1,000 $\bar{Y}(10)$, $\bar{Y}(100)$, $\bar{Y}(1,000)$ and $\bar{Y}(10,000)$ values. Your boss can see that the variability in the results you and your friends submit decreases as m increases. That is, they cluster more closely around a central value. Also, the shape of the histogram becomes more nearly symmetric and bell-shaped although it appears to be decreasing to a point.

If your boss instead plotted a histogram of the 1,000 $\sqrt{m}(\bar{Y}(m) - 2.133)$ values for each m —that is, she took each sample mean $\bar{Y}(m)$, subtracted the true steady-state mean $\mu = 2.133$, and multiplied the result by \sqrt{m} —then she would see the

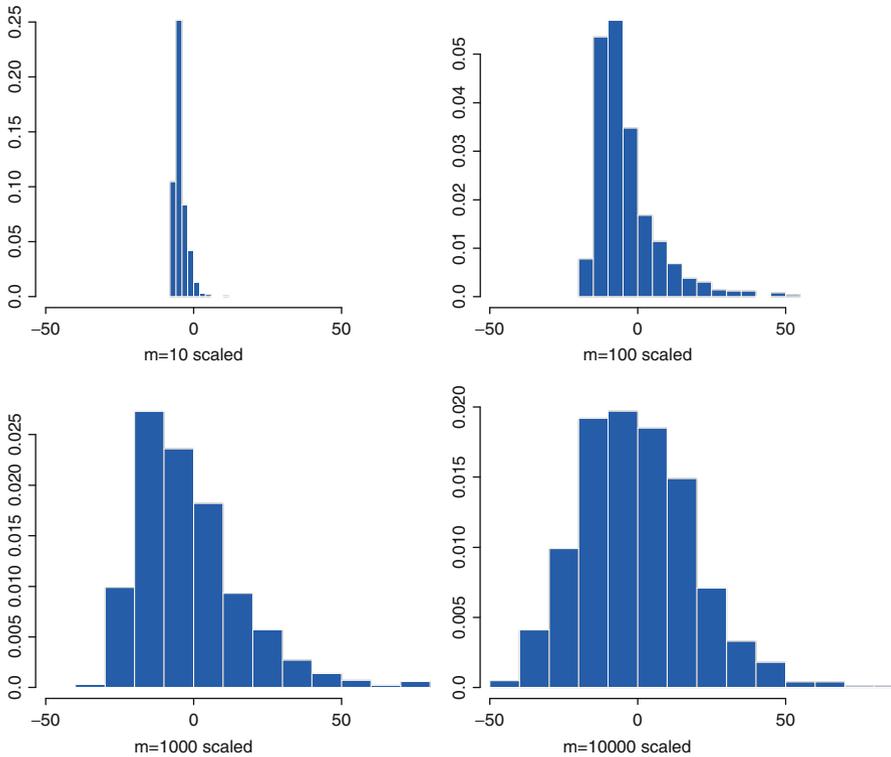


Fig. 5.3 Histogram of the 1,000 scaled and centered replication averages at $m = 10, 100, 1,000$ and $10,000$

sequence of histograms in Fig. 5.3. Notice that rather than the variability of the histograms decreasing, as happened with the raw averages, the distribution appears to be stabilizing toward a specific mean-zero (perhaps) normal distribution. This is an example of convergence in distribution as promised by the central limit theorem, which we state precisely below. We often say loosely that “the sample means become normally distributed,” but as these figures show, the scaling up by \sqrt{m} is needed to converge to a stable distribution (one whose variability is not shrinking as in Fig. 5.2).

Still looking at things from your boss’s point of view, we will summarize the sample averages in three additional ways: $\bar{Y}(m)$, the overall average of the 1,000 averages obtained at each check point; $\bar{Y}(m) - 2.133$, the difference between the overall average and the true steady-state mean; and the fraction of the 1,000 averages that differ from $\mu = 2.133$ min by more than 0.5 min. These three measures are empirical estimates of $E(\bar{Y}(m))$, $\text{Bias}(\bar{Y}(m))$ and $\text{Pr}\{|\bar{Y}(m) - 2.133| > 0.5\}$. The results are shown in Table 5.1. Notice that as m increases, the overall mean seems to be converging to 2.133, the bias seems to be going to 0, and the likelihood of a sample mean being more than 0.5 min from 2.133 seems to be going to zero. Notice

Table 5.1 Summary measures of the 1,000 sample means

m	10	100	1,000	10,000
$\bar{Y}(m)$	0.718	1.739	2.083	2.130
$\bar{Y}(m) - 2.133$	-1.415	-0.394	-0.0503	-0.003
$\frac{\#\{ \bar{Y}(m) - 2.133 > 0.5\}}{1,000}$	0.914	0.742	0.359	0.010

also that the bias is negative, a result of starting the simulation with no patients and visitors in the system so that for small m the delay in queue tends to be quite low.

Your sample means contributed to these results; they were 0.517, 2.212, 1.621, and 1.972, for $m = 10, 100, 1,000$ and 10,000, respectively, each just one of the 1,000 sample means your boss observed at each check point. None of them equal 2.133, meaning they all have error, consisting of bias and sampling variability. What does large-sample asymptotic behavior tell you about *your* single result? Knowing how things look from your boss' point of view helps provide an interpretation of the various modes of convergence as they apply to you; we define these now and relate them back to the example.

Consider a sequence of random variables $\{W_1, W_2, \dots\}$ and a constant μ . To relate the situation to the queueing problem, let $W_m = \bar{Y}(m), m = 1, 2, \dots$, be *your* sequence of sample means.

Definition 5.1. Convergence in probability: $W_m \xrightarrow{P} \mu$ if for any $\varepsilon > 0$,

$$\lim_{m \rightarrow \infty} \Pr\{|W_m - \mu| > \varepsilon\} = 0.$$

In the queueing example we noticed that as m increases, fewer and fewer of the sample means your boss saw were greater than $\varepsilon = 0.5$ away from $\mu = 2.133$. While you could be unlucky with your one result, convergence in probability says that the chances that you are unlucky become smaller and smaller. Your last sample mean, $\bar{Y}(10,000) = 1.972$, is less than 0.5 from μ .

Definition 5.2. Convergence with probability 1: $W_m \xrightarrow{a.s.} \mu$ if

$$\Pr\left\{\lim_{m \rightarrow \infty} W_m = \mu\right\} = 1.$$

The abbreviation a.s. stands for "almost surely," so convergence with probability 1 is also called almost-sure convergence.

To really make the definition of a.s. convergence complete we also need to define the underlying probability space on which the sequence W_m is defined. However, a simulation analogy captures the essence of this type of convergence: Recall that you and your friends independently and randomly selected a random-number seed for your queueing simulation. Once you pick a seed your simulation outputs are a deterministic sequence of numbers. Almost sure convergence means that with probability

1 you will select a seed such that your sequence of sample means $\bar{Y}(m)$ converges to μ in the usual sense of convergence of a sequence of numbers to a constant. Some of your friends may get there faster, and others slower, but you are all guaranteed that as m increases *your* sample mean $\bar{Y}(m)$ will converge to μ .¹

Although we have defined convergence in probability and with probability 1 as convergence of a sequence of random variables to a constant, they can also be extended to convergence to a random-variable W defined on the same probability space as W_m . However, we do not need the more general definition in this book.

To describe convergence in distribution, and specifically the central limit theorem, consider a sequence of random variables Z_1, Z_2, \dots and a random variable Z whose distribution is not a function of m . Let $F_m(z) = \Pr\{Z_m \leq z\}$ and let $F(z) = \Pr\{Z \leq z\}$. In the queueing simulation we could let $Z_m = \sqrt{m}(\bar{Y}(m) - \mu)$.

Definition 5.3. Convergence in distribution: $Z_m \xrightarrow{D} Z$ if

$$\lim_{m \rightarrow \infty} F_m(z) = F(z)$$

for all points z at which $F(z)$ is continuous. This definition extends without change to Z_m and Z being vector-valued random variables.

In the queueing simulation, the boss saw that as m got large the empirical distribution of the scaled and centered sample means became a stable, mean-zero normal distribution. Thus, F is a mean-zero normal distribution. When convergence in distribution is established to a mean-zero normal distribution it is often referred to as a “central limit theorem.”

Central limit theorems are extremely valuable. If one holds for $Z_m = \sqrt{m}(\bar{Y}(m) - \mu)$, then for large m we can say that $\sqrt{m}(\bar{Y}(m) - \mu) \approx \gamma N(0, 1)$, where $N(0, 1)$ is a standard normal random variable and γ^2 is a (usually unknown) variance constant. Or equivalently, for large m

$$\bar{Y}(m) \approx \mu + \frac{\gamma}{\sqrt{m}} N(0, 1). \tag{5.1}$$

Thus, up to the variance constant γ^2 (which can often be estimated), the central limit theorem completely characterizes the error in using $\bar{Y}(m)$ to estimate μ . Specifically, it says that the error is symmetric about μ , and normally distributed, with standard deviation γ/\sqrt{m} . This characterization permits the derivation of bounds on the error that are correct with high probability (a confidence interval). The only missing piece is the variance constant, which is why estimation of γ^2 has been such an important topic in simulation research.

¹ While this is a good analogy, it is not precisely correct since random-number generators imitate truly random numbers and in fact repeat the same (long) cycle of values over and over again.

5.2.2 Asymptotics Across Replications

The large-sample properties defined in the previous section are useful because they provide some assurance that our simulation-based estimates are converging to the performance measures of interest, and by helping to quantify the remaining error when we stop simulating. This leads to the obvious question: When can we expect such convergence properties to hold?

The most easily verifiable collection of theorems are for the i.i.d. case; in the queueing example, this is from your boss's point of view. Remember that your boss sees $n = 1,000$ (or more if you find more friends) results, each independently simulated according to the same rules. In other words, your boss sees independent and identically distributed replications.

Theorem 5.1. Weak law of large numbers (WLLN): *If Z_1, Z_2, \dots are i.i.d. with $E(Z_1) < \infty$, then*

$$\frac{1}{n} \sum_{i=1}^n Z_i \xrightarrow{P} E(Z_1).$$

Strong law of large numbers (SLLN): *If Z_1, Z_2, \dots are i.i.d. with $E(Z_1) < \infty$, then*

$$\frac{1}{n} \sum_{i=1}^n Z_i \xrightarrow{a.s.} E(Z_1).$$

Central limit theorem (CLT): *If Z_1, Z_2, \dots are i.i.d. with $E(Z_1^2) < \infty$, then*

$$\sqrt{n} \left(\frac{1}{n} \sum_{i=1}^n Z_i - E(Z_1) \right) \xrightarrow{D} \gamma N(0, 1)$$

where $\gamma^2 = \text{Var}(Z_1)$.

These results say that if your boss hired more and more friends to submit results, the average of these results would converge (both in probability and almost surely) to the mathematical expected value, and that the distribution of the (scaled) average would look more and more normally distributed.

There is a very subtle but important point to emphasize here: The centering constant in these theorems is $E(Z_1)$, which is not necessarily μ , the performance measure you are interested in estimating.² This is exactly the situation in the queueing example: For any finite m , the $E(Y(m)) \neq \mu$ (the steady-state mean) due to bias induced by the way the simulation was initialized empty and idle. As we saw in Table 5.1, results are biased low (underestimating the steady-state mean waiting time). Since we will often make many replications to control statistical error, we need to remember that replications do not affect this sort of bias. In the next section we see how this sort of bias can be tamed by increasing the length of the replication (number of customers m in the queueing example).

² Since Z_1, Z_2, \dots are i.i.d., we can select Z_1 or any of the others as the representative.

What is “ Z_1 ” in each of our other standard examples, and is $E(Z_1)$ what we want to estimate?

Stochastic activity network: Here the goal was to estimate $\theta = \Pr\{Y > t_p\}$ and our observation on replication 1 is

$$Z_1 = \begin{cases} 1, & Y_1 > t_p \\ 0, & Y_1 \leq t_p \end{cases}$$

Thus, $E(Z_1) = 1 \cdot \Pr\{Y_1 > t_p\} + 0 \cdot \Pr\{Y_1 \leq t_p\} = \theta$; so the WLLN, SLLN and CLT guarantee precisely the sort of convergence we want.

Asian option: Here the quantity of interest is

$$v = E \left[e^{-rT} (\bar{X}(T) - K)^+ \right]$$

but what we observe is

$$Z_1 = e^{-rT} \left(\widehat{\bar{X}}_1(T) - K \right)^+$$

where $\widehat{\bar{X}}_1(T)$ is a discretized version of $\bar{X}(T)$ from replication 1; discretization introduces bias. Thus, $E(Z_1) \neq v$, and no number of replications can cause $\sum_{i=1}^n Z_i/n$ to converge to v ; of course, it will converge to something not far from v if the discretization error is small.

M(t)/M/ ∞ queue: Suppose that, as in Sect. 4.2, we are interested in the maximum number of cars in the parking lot over a 24-h period. Recall that $N(t)$ was the stochastic process representing the number of cars in the lot at time t . A quantity of interest then might be

$$\mu = E \left(\max_{0 \leq t \leq 24} N(t) \right)$$

the expected value of the maximum over a 24-h period. The natural output—as illustrated in Sect. 4.2—is

$$Z_1 = \max_{0 \leq t \leq 24} N_1(t).$$

Therefore $\mu = E(Z_1)$ by definition and the convergence results in this section imply that taking more replications guides us to the right value.

The WLLN, SLLN and CLT show that sample means of i.i.d. data have very desirable properties under very general conditions. When we average across-replication outputs, these results apply. In our queueing example the replications came from you and your friends conducting independent simulations; more typically it is only you executing multiple replications with different random numbers. Assuming a good random-number generator, we treat these two ways of obtaining replications as the same. However, we have to take some care to understand the relationship between the expected value of our estimator and the expected value we are trying to estimate; when they are not the same, no quantity of replications can reduce the bias.

Remark 5.1. Convergence in probability and with probability 1 do not always hold together; if this was the case, then there would only be a need for one concept. The SLLN is called “strong” for a reason: it implies that the WLLN holds, but not the other way around. This important distinction is treated in advanced texts on probability such as Billingsley (1995).

5.2.3 Asymptotics Within Replications

Recall that in the queueing experiment you only got to see one (quite long) replication. The results that your boss saw from across many replications suggested that you might be able to assume that a strong law of large numbers or central limit theorem applies to *your* particular $\bar{Y}(m)$ as the number m of patients and visitors, rather than the number n of replications, increases. However, your data are clearly not i.i.d. since the first patient waits 0 min no matter what, and Lindley’s equation (3.3) shows that the waiting times of successive patients and visitors are dependent. Therefore, the theorems of Sect. 5.2.2 do not apply.

Unfortunately, there are no general, easily verifiable conditions for establishing convergence in the non-i.i.d. case. In fact, a rather deep understanding of the stochastic process being simulated is essential to actually prove a SLLN/WLLN or CLT for a dependent process. See Henderson (2006) and references therein for a lucid, but technically challenging exploration of this point.

That said, the desired large-sample properties often hold. Typically heuristic justifications are offered based on analogies with simpler processes that have been studied: Stationary Markovian queues where the service capacity exceeds the customer load; stationary, irreducible, positively recurrent Markov chains; time series processes such as the AR(1); and regenerative processes are a few examples. The Appendix to this chapter discusses regenerative structure in more detail; interested readers can refer to Haas (2002) and Glynn (2006).

Roughly, conditions such as the following must hold:

- Neither the logic of the system nor the input processes that drive the simulation can change over time, and there can be no periodic (cyclic) behavior.
- The state space of the simulation may not decompose into distinct subsets such that the simulation gets trapped in one or the others based on the particular random numbers used.
- The state of the simulation in the distant future must be effectively independent of the state in the past. In particular, the initial state at the start of a simulation replication must become irrelevant when simulating far enough into the future. This requires at least a minimal level of randomness and forgetfulness.

We now suppose that these conditions are satisfied, and the output process Y_1, Y_2, \dots within a replication is such that $Y_m \xrightarrow{D} Y$, a random variable Y whose distribution does not depend on the time index. We refer to Y as the steady state, and the existence of Y implies that limit properties such as $\mu = E(Y)$ are well defined.

Even though Y_m has a steady state, the estimator $\bar{Y}(m)$ may still be biased (as in the $M/G/1$ queue). If we simulate long enough, then we want the bias in $\bar{Y}(m)$ to disappear. Define the *asymptotic bias* as

$$\begin{aligned}\beta &= \lim_{m \rightarrow \infty} m(\mathbb{E}(\bar{Y}(m)) - \mu) \\ &= \lim_{m \rightarrow \infty} m \left(\mathbb{E} \left(\frac{1}{m} \sum_{i=1}^m Y_i \right) - \mu \right) \\ &= \sum_{i=1}^{\infty} (\mathbb{E}(Y_i) - \mu)\end{aligned}\tag{5.2}$$

(Whitt 2006). For the bias in $\bar{Y}(m)$ to disappear, the bias of the individual outputs Y_i must fall off fast enough so that $\beta < \infty$; otherwise we cannot overcome the bias by increasing m . Thus, for large m , $\text{Bias}(\bar{Y}(m)) \approx \beta/m$. Estimating β is very difficult, so the fact that it must fall off quickly (as $1/m$) is important because we can overwhelm it with a long replication.

For any random variables Y_1, Y_2, \dots, Y_m ,

$$\text{Var}(\bar{Y}(m)) = \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m \text{Cov}(Y_i, Y_j).\tag{5.3}$$

This formula is a direct extension of the basic result $\text{Var}(Y_1 + Y_2) = \text{Var}(Y_1) + \text{Var}(Y_2) + 2\text{Cov}(Y_1, Y_2) = \text{Cov}(Y_1, Y_1) + \text{Cov}(Y_2, Y_2) + \text{Cov}(Y_1, Y_2) + \text{Cov}(Y_2, Y_1)$. Notice that the expression has m^2 terms and we only have m observations, so there is no way to estimate it directly. What we hope, and can sometimes prove, is that the dependence structure of Y_1, Y_2, \dots also stabilizes for large m , not just the marginal properties such as mean and variance. The standard assumption is that beyond some point we can treat the process as being *covariance stationary*, meaning that for m large enough, $\sigma^2 = \text{Var}(Y_m)$ and $\rho_k = \text{Corr}(Y_m, Y_{m+k})$ are no longer a function of m . The quantity ρ_k is called the *autocorrelation* at lag k . For a covariance stationary process (5.3) simplifies to

$$\text{Var}(\bar{Y}(m)) = \frac{\sigma^2}{m} \left(1 + 2 \sum_{k=1}^{m-1} \left(1 - \frac{k}{m} \right) \rho_k \right).\tag{5.4}$$

Therefore, for a covariance stationary process the *asymptotic variance* is

$$\begin{aligned}\gamma^2 &= \lim_{m \rightarrow \infty} m \text{Var}(\bar{Y}(m)) = \lim_{m \rightarrow \infty} \sigma^2 \left(1 + 2 \sum_{k=1}^{m-1} \left(1 - \frac{k}{m} \right) \rho_k \right) \\ &= \sigma^2 \left(1 + 2 \sum_{k=1}^{\infty} \rho_k \right).\end{aligned}\tag{5.5}$$

For the $\text{Var}(\bar{Y}(m))$ to go to 0 with m we need $\gamma^2 < \infty$. Covariance stationarity alone does not guarantee this, since Y_i having “heavy tails” ($\sigma^2 = \infty$) or long-range dependence (ρ_k not falling off fast enough) could cause γ^2 to be infinite. When γ^2

exists then $\text{Var}(\bar{Y}(m)) \approx \gamma^2/m$ for large m . Further, if ρ_k falls off quickly relative to m then observations far enough apart are effectively uncorrelated and there is some hope of estimating γ^2 (see Chap. 8).

Bias is a systematic error, while variance reflects error due to random sampling. It might appear that they are equally important, because they both decrease as $1/m$. However, variance is not a measure of error, but of squared error: $\text{Var}(\bar{Y}(m)) = \text{E} \left[\{\bar{Y}(m) - \text{E}(\bar{Y}(m))\}^2 \right]$. A more relevant measure is the standard error, $\sqrt{\text{Var}(\bar{Y}(m))}$, which decreases as $1/\sqrt{m}$. Thus, even though both are important, in a large-sample sense sampling variability dominates bias. A combined measure that treats bias and variance equally is the *mean squared error*,

$$\begin{aligned} \text{MSE}(\bar{Y}(m)) &= \text{Bias}^2(\bar{Y}(m)) + \text{Var}(\bar{Y}(m)) \\ &\approx \frac{\beta^2}{m^2} + \frac{\gamma^2}{m}. \end{aligned} \tag{5.6}$$

Remark 5.2. Although we defined modes of convergence in terms of a discrete-time process Z_1, Z_2, \dots , there are completely analogous definitions for continuous-time processes, $\{Z(t), t \geq 0\}$ (e.g., the queue-length process $Y(t)$ instead of waiting time in the queue Y_1, Y_2, \dots). In particular, the CLT for a continuous-time process shows that

$$\lim_{t \rightarrow \infty} \sqrt{t} \left(\frac{1}{t} \int_0^t Z(t) dt - \mu \right) \xrightarrow{D} \gamma \text{N}(0, 1).$$

5.2.4 Beyond Sample Means

Not every performance measure is estimated by a sample mean (e.g., the variance). However, many estimators can be thought of as functions of sample means. A few key results are helpful for understanding the simulation literature.

Theorem 5.2 (Converging together lemma). *Suppose $Z_m \xrightarrow{D} Z$ and $W_m \xrightarrow{P} \eta$, where η is a constant. Then*

$$(Z_m, W_m) \xrightarrow{D} (Z, \eta).$$

This theorem states that if Z_m converges in distribution to Z , while W_m converges in probability to a constant, then no matter what their joint distribution is, they converge together to these limits. The converging together lemma is often paired with the following:

Theorem 5.3 (Continuous mapping theorem). *Suppose Z_m and Z are random variables in \mathfrak{R}^d ; that $Z_m \xrightarrow{D} Z$; and that $h(\cdot)$ is a continuous function from \mathfrak{R}^d to \mathfrak{R} . Then*

$$h(Z_m) \xrightarrow{D} h(Z).$$

This result states that when random variables converge in distribution, then functions of them may also converge. The statement of the theorem is actually stronger than necessary: h may have discontinuities in a set D as long as $\Pr\{Z \in D\} = 0$.

These two results turn the CLT into a practically useful tool. Suppose we have an estimator Z_n that satisfies the CLT

$$\sqrt{n}(Z_n - \mu) \xrightarrow{D} \gamma N(0, 1)$$

and we have an estimator $S_n^2 \xrightarrow{a.s.} \gamma^2$ (remember that almost sure convergence implies convergence in probability). Then if we let $h(a, b) = a/\sqrt{b}$, the converging together lemma and the continuous mapping theorem imply that

$$h(\sqrt{n}(Z_n - \mu), S_n^2) = \frac{\sqrt{n}(Z_n - \mu)}{S_n} \xrightarrow{D} \frac{\gamma N(0, 1)}{\gamma} = N(0, 1).$$

Thus, for large n we can argue that $Z_n \approx \mu + N(0, 1) S_n / \sqrt{n}$. This provides an asymptotic justification for the usual confidence interval $Z_n \pm 1.96 S_n / \sqrt{n}$, where 1.96 is the 0.975 quantile of the standard normal distribution. More generally, the converging together lemma and continuous mapping theorem are useful for establishing the validity of “plug-in estimators” where unknown constants (in this case the variance constant γ^2) are replaced with convergent estimators of them.

Here is another application of the continuous mapping theorem: Suppose that Z_1, Z_2, \dots are i.i.d. with finite mean μ and variance σ^2 , and we want to estimate $h(\mu)$ for continuous and twice differentiable function h .

The continuous mapping theorem implies that $h(\bar{Z}(n))$ converges to $h(\mu)$. Since we will have to stop our simulation short of ∞ , what can we say about the estimator $h(\bar{Z}(n))$? Using a Taylor series expansion

$$h(\bar{Z}(n)) \approx h(\mu) + h'(\mu)(\bar{Z}(n) - \mu)$$

so that provided $h'(\mu) \neq 0$

$$\begin{aligned} \sqrt{n}(h(\bar{Z}(n)) - h(\mu)) &\approx h'(\mu)\sqrt{n}(\bar{Z}(n) - \mu) \\ &\xrightarrow{D} h'(\mu)\sigma N(0, 1) \end{aligned}$$

by the CLT and continuous mapping theorem. Therefore, $\text{Var}(h(\bar{Z}(n))) \approx h'(\mu)^2 \sigma^2 / n$ for large n . Taking the Taylor series out to one more term, a similar argument shows that

$$\text{Bias}(h(\bar{Z}(n))) = \mathbb{E}(h(\bar{Z}(n))) - h(\mu) \approx \frac{1}{2} \frac{h''(\mu)\sigma^2}{n}. \quad (5.7)$$

This approach for deriving the variance and bias of one statistic that is a function of another is usually called the “Delta method.” The base statistic need not be a sample mean, but it does need to be a statistic whose asymptotic properties are known.

Appendix: Regenerative Processes and Steady State

How do we establish that a simulated system has a steady state? Or that simulation-based estimators satisfy a SLLN or CLT? As indicated earlier, this is a difficult question in general. In some cases we can argue that the simulation generating the output is a *regenerative process*. For deeper discussion beyond this appendix, see Haas (2002) and Glynn (2006).

Our goal is to prove that the output process Y_t has a steady state, and that the sample mean satisfies a SLLN and CLT. We will handle the continuous-time ($t \geq 0$) and discrete-time ($t = 0, 1, \dots$) cases together.

Suppose that in the same simulation we can identify a renewal process $\{S_i, i = 0, 1, 2, \dots\}$; by a renewal process we mean that

$$\begin{aligned} S_0 &= 0 \\ S_i &= A_1 + A_2 + \dots + A_i \end{aligned}$$

where A_i are i.i.d. with $\Pr\{A_i = 0\} < 1$ and $\Pr\{A_i < \infty\} = 1$. The output processes $\{Y_t; t \geq 0\}$ or $\{Y_t; t = 0, 1, \dots\}$ are *regenerative processes* if taken together

$$\{Y_t; S_i \leq t < S_{i+1}\}, i = 0, 1, 2, \dots$$

are i.i.d. Notice that the output process Y_t may be highly dependent; being regenerative means that the behavior of $\{Y_t; S_i \leq t < S_{i+1}\}$ and $\{Y_t; S_{i+1} \leq t < S_{i+2}\}$, say, are independent of each other, and have the same probability distribution. For this reason, regenerative processes are said to probabilistically “start over” at the renewal times $t = S_0, S_1, S_2, \dots$.

For instance, an ergodic discrete or continuous-time Markov chain starts over every time it enters a fixed state, due to the Markov property. Many queueing systems start over probabilistically each time a customer arrives to find the system empty and idle.

Let

$$Z_i = \int_{S_{i-1}}^{S_i} Y_t dt$$

be the accumulated output over the i th renewal cycle. The integral becomes a sum for discrete-time output processes. If Y_t is regenerative (with right-continuous sample paths if S_i is continuous valued), $E(|Z_1|) < \infty$, $E(A_1) < \infty$, and A_1 is aperiodic,³ then the following hold:

- $Y_t \xrightarrow{D} Y$ (the process has a steady-state distribution);
- $\lim_{t \rightarrow \infty} t^{-1} \int_0^t Y_s ds \xrightarrow{a.s.} E(Z_1)/E(A_1)$ (the sample mean converges to the expected total output per cycle divided by the expected length of a cycle); and

³ Aperiodic means that for A_1 continuous valued, there is no $d > 0$ such that $\sum_{n=0}^{\infty} \Pr\{A_1 = nd\} = 1$; and for A_1 discrete valued, there is no integer $d \geq 2$ such that $\sum_{n=0}^{\infty} \Pr\{A_1 = nd\} = 1$.

- $\mu = E(Y) = E(Z_1)/E(A_1)$ (the steady-state mean is the almost sure limit of the sample mean).

Thus, regenerative structure establishes that there is a steady-state distribution and that as the length of the replication increases the sample mean satisfies a SLLN for the steady-state mean.

Now let $V_i = Z_i - \mu A_i$; this is the cumulative output from the i th cycle minus the predicted cumulative output based only on the length of the cycle. If it is also the case that $E(V_1^2) < \infty$, then

- $\lim_{t \rightarrow \infty} \sqrt{t}(\bar{Y}_t - \mu) \xrightarrow{D} \gamma N(0, 1)$, where $\bar{Y}_t = t^{-1} \int_0^t Y_s ds$; and
- The asymptotic variance $\gamma^2 = E(V_1^2)/E(A_1)$.

Therefore, the sample mean satisfies a CLT, and the asymptotic variance can be expressed in terms of properties of a regenerative cycle.

Establishing the existence of steady state using the regenerative argument amounts to identifying the renewal process S_i and arguing that $\{Y_t; S_i \leq t < S_{i+1}\}$ are i.i.d. Typically S_1, S_2, \dots are times such that the state of the simulation is identical, and the probability distributions of all pending events are also the same. Further, one needs to establish that $\Pr\{A_i < \infty\} = 1$; that is, the time between regenerations is finite.

Consider the $M/G/1$ queue of Sect. 5.2.1. Let S_i be the i th time an arriving customer finds the system completely empty. At such times nothing that has happened in the past is relevant: there are no other customers in the system, the distribution of the time until the next arrival is exponential with parameter $\lambda = 1$, and the service time of the customer that just arrived is Erlang with mean $\tau = 0.8$ and three phases. The time between such occurrences is finite because $\lambda \tau < 1$, meaning that the system can keep up and therefore will be empty from time to time.

Exercises

1. Suppose that Z_1, Z_2, \dots, Z_n are i.i.d. with finite mean μ and variance σ^2 . Use the WLLN and continuous mapping theorem to show that $S^2 \xrightarrow{P} \sigma^2$.
2. Look up and write down corresponding convergence definitions for continuous-time processes.
3. Suppose that Z_1, Z_2, \dots, Z_n are i.i.d. with finite mean μ , and variance σ^2 , and that $h(\cdot)$ is twice continuously differentiable at μ . Show that

$$n[E(h(\bar{Z}(n))) - h(\mu)] \longrightarrow \frac{1}{2}h''(\mu)\sigma^2.$$

4. Derive Eq. (5.3).

5. Starting with Eq. (5.3), show that for a covariance stationary process

$$\text{Var}(\bar{Y}(m)) = \frac{\sigma^2}{m} \left(1 + 2 \sum_{k=1}^{m-1} \left(1 - \frac{k}{m} \right) \rho_k \right).$$

6. Derive Eq. (5.7).
7. To prove that the results obtained by the Delta method are correct, conditions are needed on the higher-order terms in the Taylor series expansion. Derive or research these conditions.
8. In Sect. 5.2.4 it was stated that “The continuous mapping theorem implies that $h(\bar{Z}(n))$ converges to $h(\mu)$.” What mode of convergence does it satisfy? *Hint:* Convergence in distribution to a constant implies convergence in probability to that constant.
9. For the AR(1) process described in Sect. 3.3, derive the asymptotic MSE of the sample mean $\bar{Y}(m)$.
10. For the following scenarios, state whether the estimator will be biased and explain why or why not.
 - a. For the stochastic activity network, we want to estimate the mean time to complete the project, and for our estimator we use Eq. (3.11).
 - b. A European option is based on $X(T)$, the value of the asset at time T . To estimate the expected value of a European option and we use $e^{-rT}(X(T) - K)^+$.
 - c. For an $M/G/1$ queue, we are interested in estimated the mean waiting time of the tenth customer to arrive when the system starts empty, and we use Y_{10} from Eq. (3.3) as our estimator.
 - d. Suppose that the parking lot represented by the $M(t)/M/\infty$ queue is open from 8 a.m. to 11 p.m. daily. A replication is defined by 1 day, and the garage starts empty each day. We want to estimate the expected value of the number of hours in a day when the garage has more that 2,000 cars in it. We take as our output

$$Z = \int_0^T M(t) dt$$

where $T = 15$ h and

$$M(t) = \begin{cases} 1, & N(t) > 2,000 \\ 0, & N(t) \leq 2,000. \end{cases}$$

11. In addition to the AR(1), another surrogate model that is used to represent steady-state simulation output is the MA(1):

$$Y_i = \mu + \theta X_{i-1} + X_i, \quad i = 1, 2, \dots, m$$

where X_1, X_2, \dots are i.i.d. $(0, \sigma^2)$ random variables, X_0 may be random or fixed, and $|\theta| < 1$. Show for this process that $\text{Var}(Y_i) = (1 + \theta^2)\sigma^2$, and

$$\text{Cov}(Y_i, Y_{i+j}) = \begin{cases} \theta \sigma^2, & j = 1 \\ 0, & j > 1. \end{cases}$$

Use this information to derive the asymptotic MSE of $\bar{Y}(m) = \sum_{i=1}^m Y_i/m$ as an estimator of μ .

12. A mode of convergence not covered in this chapter is *convergence in mean square* (also called mean-square convergence). Look up the definition and describe how it relates to the other modes of convergence that were covered.

Chapter 6

Simulation Input

This chapter covers simulation input, which includes the following:

Input modeling: Selecting (and perhaps fitting) the probability models that represent the uncertainty in the simulated system.

Random-variate generation: Representing the inputs as transformations of i.i.d. $U(0, 1)$ random variables.

Random-number generation: Producing a good approximation to realizations of i.i.d. $U(0, 1)$ random variates.

We present the topics in this order because input modeling is the objective, while random-variate generation and random-number generation are means to an end.

6.1 Input Modeling Overview

In many cases, the inputs themselves are not of great interest; what is of interest are the outputs that the inputs imply. For random-variate generation we can use provably correct methods; for random-number generation we can use generators with well-documented properties. Input modeling, on the other hand, may require judgement and is therefore open to more error. Since fitting distributions to data using methods such as maximum likelihood estimation (MLE) is a standard topic in statistics, our focus will be more on avoiding errors and less on the mechanics of fitting. We also discuss selecting input models when there are no data.

6.1.1 An Input Modeling Story

To motivate the important topics in input modeling, we will walk through a generic input modeling story, pausing from time to time to discuss it. The story represents a typical input modeling approach that makes standard assumptions; the discussion

outlines pitfalls and issues. The story parallels the simulation framework described in Sect. 5.1: There is a conceptual system of interest that does not yet exist, but there is a corresponding real system that is similar enough that some input data and system logic can be observed. The goal is to build a simulation of the conceptual system design.

Recall the hospital reception problem, Example 3.2: Under consideration was replacing the human receptionist by an electronic kiosk with a touch screen that might be a bit slower, or more variable, for patients and visitors who are less comfortable interacting with a touch screen. Therefore, the hospital management engineers wanted to evaluate how much additional delay that this might cause.

Hospital records were available to estimate the overall arrival rate. The physical basis for the Poisson distribution—arrivals from a large population of potential customers making independent decisions about when to arrive—suggests using it as the arrival-process input model. (For this reason exponentially distributed interarrival times are the default in some commercial simulation languages.)

To “fit” this input model, hospital records were used to count the number of arrivals over some time interval $[0, T]$ and the arrival rate λ was estimated by

$$\hat{\lambda} = \frac{N(T)}{T}$$

where $N(t)$ is the number of arrivals that occurred by time t . This turns out to be an unbiased, maximum likelihood estimator of λ if the arrival process is really Poisson.

Notice that we are not just treating the arrival process as Poisson, but as *stationary* Poisson (constant arrival rate), which is a significantly stronger approximation. Unfortunately, if the arrival process was nonstationary Poisson, there is nothing about the estimator $\hat{\lambda}$ to alert us; when we carry out this calculation we will get a number, even if the arrival rate is changing radically throughout the day as in the parking lot Example 3.1. *The only way to detect nonstationarity is to look at detailed arrival data, either the counts over smaller time intervals, or the interarrival times themselves if available.*

Suppose that the arrival rate is indeed constant throughout the day. Let T be the length of 1 day, and suppose we had recorded patient arrivals counts for, say, m days. Let $N_i(T), i = 1, 2, \dots, m$ be the counts from the m days. Then a statistical test could be used to see if the data support the hypothesis that the counts are Poisson. “Goodness-of-fit tests” are formulated as

H_0 : the chosen distribution with the estimated parameters is correct

H_1 : the chosen distribution with the estimated parameters is incorrect.

There are at least three reasons why the test might reject the Poisson input model:

1. The distribution of the data is substantially different from Poisson. For instance, if the reception desk only serves patients with appointments, and a fixed number

of appointments are scheduled each day, then the counts are nearly identical and therefore not Poisson. Also, scheduled appointments are typically spread evenly throughout the day and therefore arrivals are not as random as a Poisson model would predict.

2. The distribution of arrivals might differ by day of the week. For instance, Saturdays could be especially busy. Saturday arrival counts, considered by themselves, might be well-modeled as stationary Poisson, but not all days taken together. Thus, even if arrivals throughout a day can be treated as stationary Poisson there could still be nonstationarity across days.
3. We might have too much data. Why could this cause the hypothesis to be rejected? Treating the counts as Poisson is always an approximation if we are dealing with a physical (as opposed to simulated) system; real data do not come from probability distributions. Thus, any valid goodness-of-fit test will, if given enough data, reject every distribution. This does not mean that goodness-of-fit tests are useless, but it does mean that they should be used with caution. Goodness-of-fit tests make the most sense when there is a physical basis for the distribution choice and we want a warning if the data vary substantially from that choice.

Hospital records would certainly provide counts. But if arrival times themselves are logged, then they would provide more detail, including interarrival times. If the arrival process is (stationary) Poisson, then the distribution of the interarrival times is exponential, and this fact provides another way to validate the input model.

The management engineers collected data on people interacting with the kiosk from a trial study with the vendor. For the service time distribution there is no strong process physics guiding the choice of distribution, as there was for the arrival process. Therefore, the management engineers analyzed the data using distribution fitting software and chose the distribution that the software recommends as the “best fit.”

Distribution fitting software typically fits every relevant distribution it can using methods such as MLE, and then it recommends a candidate using some kind of score, often the best (smallest) goodness-of-fit statistic. Fitting all relevant distributions makes comparisons among the choices convenient, but otherwise this is a questionable approach. First, any score summarizes all aspects of the fit by a single number, without regard to where the lack of fit occurs, or if it matters. Second, if a goodness-of-fit test is to have any meaning at all, then there must be a distribution that was hypothesized due to some additional information (as we did with the Poisson choice for the arrival process); *there is no statistical meaning that can be attached to the distribution having the smallest test score from among an arbitrary collection.*

Are the data collected during the vendor trial representative of what will occur in the hospital? If the trial “patients and visitors” are vendor or hospital staff, then maybe not. Is learning anticipated? That is, are patients and visitors expected to get better at using the kiosk over time? If so, then any distribution fit to the trial data

will not represent long-run behavior. *The message is that data have to be relevant to the situation at hand to be useful.*

Recall that the Pollaczek–Khinchine formula (Eq. (3.4)) gives the steady-state expected waiting time in an $M/G/1$ queue as

$$E(Y) = \frac{\lambda(\sigma^2 + \tau^2)}{2(1 - \lambda\tau)}$$

where λ is the arrival rate of the Poisson arrival process and (τ, σ^2) are the mean and variance of the service-time distribution. One outcome of input modeling in this case is to establish values $\hat{\lambda}$, $\hat{\tau}$ and $\hat{\sigma}^2$ for these parameters. Because this is an $M/G/1$ queue, it allows us to make a few observations:

- In this situation, at least in steady state, the particular service-time distribution chosen by the software does not matter *as long as it gets the mean and variance right*. Fortunately, it is often true—which is not the same as always true—that simulation results are not overly sensitive to the specific input distributions provided key characteristics are correct. In fact, we rely on this because *any* distribution we choose is an approximation of the physical process; there is no true distribution.
- Since our parameters are estimates, they will almost certainly be wrong. That is, $\hat{\lambda} \neq \lambda$, $\hat{\tau} \neq \tau$, and $\hat{\sigma}^2 \neq \sigma^2$. While the expected waiting time is insensitive to the choice of service-time distribution, it is sensitive to the values of the mean and variance, τ and σ^2 . Further, if $\hat{\lambda}\hat{\tau} > 1$, then there is no steady-state distribution of waiting time because the queue cannot keep up. While in reality the queue might be able to keep up, a sample of data could yield $\hat{\lambda}\hat{\tau} > 1$, particularly when $\hat{\lambda}$ comes from one source (e.g., the hospital records), and $\hat{\tau}$ comes from another (the vendor study).
- Although the distribution is not critical for the service process, it is for the arrival process: that is, the Pollaczek–Khinchine result does not depend on the particular service-time distribution except through its mean and variance, but it does depend on the arrival process being Poisson. If the arrival process is more or less variable than Poisson, then the congestion will be, respectively, more or less than what is predicted by the $M/G/1$ model.

What this “view through the queue” emphasizes is that the goal of input modeling is to get relevant outputs. This frees us from a futile search for the “true, correct distribution” and toward trying, as best we can, to approximate the information and data we have about the input processes.

6.1.2 Characteristics of Input Processes

It is difficult to formally define which random variables in a stochastic simulation are inputs, as opposed to outputs or intermediate random variables between inputs

and outputs. Informally, inputs are stochastic processes whose (joint) distributions are treated as “known” in the simulation. Typically, inputs represent the uncertainty in the real world at its most basic level. We want inputs to have characteristics that match the real world, but we are usually not interested in those characteristics in and of themselves. As noted above, we want input models that provide relevant outputs.

We will divide input models into two classes:

Univariate input models: A univariate input random variable X is fully specified by its marginal distribution, F_X ; when we need multiple realizations X_1, X_2, \dots then they are i.i.d. Examples of univariate inputs include the times to failure of several electronic components of the same type; the bed occupancy times of patients in a hospital; and the number of fans who arrive to a basketball game. We often model univariate inputs using parametric probability distributions, such as the Poisson, exponential, lognormal and Weibull, with parameters that need to be tuned to the situation at hand. A particular parametric distribution may be justified by the physical basis of the input process or because it seems to provide a good fit to real-world data. We discuss selecting and fitting parametric input models in Sect. 6.2. Alternatively, we may simply reuse the data themselves; this is described in Sect. 6.2.4. When we do not have any relevant data, then other means for specifying F_X are required; see Sect. 6.2.6.

Multivariate input models: Multivariate input models define collections of random variables that are in some way related, as opposed to being i.i.d. realizations from a common distribution F_X . Examples of multivariate inputs include the quantity of whole milk, 2% milk and skim milk ordered by a grocery store; the times to failure of four tires on the same car; the arrival times of customer calls to a call center; and the returns on each of the bonds held in a portfolio. A multivariate input model must directly or indirectly specify the marginal distribution of each component random variable, as well as the dependence among them. In principle, we could define the full joint probability distribution of the multivariate input, but this may be difficult in practice for the following reasons:

- The multivariate input is a (conceptually) infinite sequence, S_1, S_2, S_3, \dots , which means that we cannot write a full joint distribution and need some other way to express how this sequence is constructed.
- We do not know the full joint distribution of, say (X_1, X_2, X_3) , but we do know the marginals and have a partial characterization of the dependence among them, such as $\text{Corr}(X_i, X_j)$ for all $i \neq j$.
- The marginal distributions (and perhaps the dependence as well) are a function of time; we refer such input processes as being *nonstationary*.

Multivariate input models for (possibly nonstationary) arrivals and for random vectors are described in Sect. 6.3.

6.2 Univariate Input Models

The situation addressed in this section often occurs in practice: We have a sample of real-world data, denoted X_1, X_2, \dots, X_m , that we are willing to model as being i.i.d observations with a stable (unchanging) distribution F_X , and we want to generate random variates from F_X to drive our simulation. A standard approach is to somehow choose a parametric family of distributions to represent F_X and to fit the distribution to the data by estimating values of its parameters. Let $F(x; \theta)$ denote the parametric family with parameter vector θ . For instance, $F(x; \theta)$ might be the lognormal family with parameters $\theta = (\mu, \sigma^2)$, its mean and variance. Then “fitting” means using some appropriate estimator $\hat{\theta} = \hat{\theta}(X_1, X_2, \dots, X_m)$, that is a function of the sample of real-world data, as the choice for θ .

The fact that we will use $F(x; \hat{\theta})$ for variate generation is important in terms of how we think about evaluating the “goodness of fit.” Let \hat{X} denote a random variable with distribution $F(x; \hat{\theta})$. To be clear on this point, \hat{X} is not one of the real-world data values X_1, X_2, \dots, X_m ; instead, it is a random variate generated from the distribution $F(x; \hat{\theta})$, where the parameters $\hat{\theta}$ are functions of the real-world data.

What makes a good fit? There are two paradigms:

Inference: The classical paradigm is to hypothesize that there is a true but unknown distribution F_X , and it is a member of a parametric family $F(x; \theta)$. In this paradigm we try to discover the true F_X using a method that has good properties, averaged over the possible distributions and data we could see. Maximum likelihood estimators of θ arise from this paradigm, as does testing of the hypothesis $H_0: F(x; \hat{\theta}) = F_X$. *In brief, the classical approach treats input modeling as a problem of statistical inference about the true distribution F_X , and employs methods that have good statistical properties.*

Matching: The matching paradigm is concerned with how closely the properties of \hat{X} match those of the real-world sample of data X_1, X_2, \dots, X_m . For instance, it might be desirable to have

$$E(\hat{X} | X_1, X_2, \dots, X_m) = \bar{X}.$$

Here the expectation is with respect to $F(x; \hat{\theta})$. This condition indicates that the expected value of a random variable with distribution $F(x; \hat{\theta})$ is the same as the sample mean of the data that was used to fit it. Depending on the application, there are many other properties of the data that it might be important to match, for instance certain tail percentiles. *In brief, the matching paradigm treats input modeling as capturing key features of the real-world data X_1, X_2, \dots, X_m .*

In practice, univariate input modeling often combines aspects of the inference and matching approaches, so we will consider the balance between the two. We do not provide details about any estimation methods (such as maximum likelihood) or hypothesis tests (such as the Kolmogorov–Smirnov test) as these are well covered in

basic statistics texts or other simulation references, such as Law (2007); also, they are implemented in input modeling software. Instead, we provide insight for proper use.

6.2.1 Inference About Univariate Distributions

Many of the well-known parametric families arise from particular process physics, often in some sort of limit. Perhaps the most well-known example is the normal distribution: if the random variable X is obtained by the summation of a large number of component random variables, then the distribution of X might tend to be normal by the central limit theorem. See Sect. 5.2.2 for the central limit theorem for the sum of i.i.d. random variables; there are also versions that apply to sums of non-identical and even dependent random variables (see for instance Lehmann 2010).

For example, consider a worker who assembles laptop computers by hand. There are many steps in the assembly process, each of which may be modeled as a random variable. Because the total time to assemble a laptop is the sum of all of these component times, the nature of the process suggests that the total time to assemble a laptop might be well-modeled by a normal distribution.

Many discrete distributions are naturally defined by the process physics that they represent: For instance, the binomial distribution describes the number of successes in a fixed number of independent and identically distributed trials, while the negative binomial distribution describes the number of such trials required to achieve a given number of successes. Input modeling means deciding what constitutes a “trial” and a “success.”

The inference paradigm is justified by having strong process physics that supports the choice of parametric family. In this section, we describe the theoretical support for selecting some common distributions (Poisson, lognormal, Weibull and gamma) as input models. A comprehensive source for this type of insight is the series of books by Johnson, Kotz and co-authors (Johnson et al. 1994, 1995, 1997, 2005; Kotz et al. 2000).

6.2.1.1 The Poisson Distribution

An arrival process resulting from a large population of potential arrivals who make independent, but infrequent, decisions about when to arrive tends to be Poisson.

In a *renewal process*, the interarrival times between successive customers or entities, denoted A_1, A_2, \dots , are i.i.d. nonnegative random variables with distribution G . Two stochastic processes that are derived from the A_n are

$$S_n = \begin{cases} 0, & n = 0 \\ \sum_{i=1}^n A_i, & n = 1, 2, \dots \end{cases}$$

$$N(t) = \max\{n \geq 0 : S_n \leq t\}.$$

Thus, S_n is the time of the n th arrival, while $N(t)$ is the number of arrivals by time $t \geq 0$. The continuous-time process N is called the *arrival-counting process*.

Now consider a collection of independent renewal processes, with their associated arrival-counting processes

$$N_1(t), N_2(t), \dots, N_h(t)$$

and each with their own interarrival-time distribution $G_i, i = 1, 2, \dots, h$. Think of these as independent sources of arrivals, maybe different types of customers, or customers from different locations. The overall number of arrivals is the superposition arrival-counting process

$$C_h(t) = \sum_{i=1}^h N_i(t).$$

This is just the total number of arrivals from all sources by time t .

Now suppose that the number of independent arrival processes h is increasing, and consider

$$C_\infty(t) = \lim_{h \rightarrow \infty} \sum_{i=1}^h N_i(t).$$

We would expect this to blow up, but we impose two conditions on the interarrival distributions G_i as h increases:

$$\lim_{h \rightarrow \infty} \max_{i=1,2,\dots,h} G_i(t) = 0 \tag{6.1}$$

$$\lim_{h \rightarrow \infty} \sum_{i=1}^h G_i(t) = \lambda t. \tag{6.2}$$

Recall that $G_i(t)$ is the probability that an interarrival time is less than or equal to t for the i th renewal arrival process; therefore, Condition (6.1) can be interpreted as implying that as we superpose more and more processes, the times between arrivals for each individual process are tending to become longer and longer. Stated differently, when the number of arrival sources h is large, the arrivals from each source are rare.

To interpret the second condition, let A_{1i} be the interarrival time from 0 until the first arrival for process i . Suppose that t is small, so that there is little chance of more than 1 arrival from each source. Then

$$E[C_h(t)] \approx E \left[\sum_{i=1}^h I(A_{1i} \leq t) \right] = \sum_{i=1}^h G_i(t).$$

So Condition (6.2) can be interpreted as implying that the limiting arrival process has a stable arrival rate.

The following result shows that under these conditions, $C_h(t)$ will have a Poisson distribution as $h \rightarrow \infty$; that is, as we superpose more and more renewal arrival processes, the overall arrival process tends to be Poisson.

Theorem 6.1 (Poisson as Superposition of Renewal Processes). *If Condition (6.1) holds, then*

$$\lim_{h \rightarrow \infty} \Pr\{C_h(t) = j\} = \frac{e^{-\lambda t} (\lambda t)^j}{j!}, \quad j = 0, 1, \dots$$

if and only if Condition (6.2) holds.

A proof of an even more general version of this result can be found in Karlin and Taylor (1975).

Consider the people living in the Deerfield, Highland Park, and Northbrook communities north of Chicago who might call the local Subaru dealer to make an appointment for repair. Each owner's calls are infrequent. In addition, there is no reason to coordinate with other Subaru owners about when to call for service; thus, their calls are independent. The theorem suggests that the number of calls for appointments over any time period might be modeled as having a Poisson distribution. Notice, however, that if the Subaru dealer schedules a fixed number of repairs per day, then the number of actual arrivals (as opposed to calls) to the dealer will not be Poisson since a bound on appointments imposes dependence on the arrival process.

6.2.1.2 The Lognormal Distribution

A random variable formed as the product (multiplication) of independent, positive random variables tends to have a lognormal distribution.

By definition, if a random variable X is such that $\ln(X)$ has a normal distribution, then X has a lognormal distribution. This connection gives rise to one asymptotic justification for selecting the lognormal distribution as an input model: Suppose that Z_1, Z_2, \dots are i.i.d. positive random variables, and

$$X = \prod_{i=1}^m Z_i.$$

Then

$$W = \ln(X) = \sum_{i=1}^m \ln(Z_i).$$

Thus, if $E(\ln(Z_1)^2) < \infty$, then the central limit theorem establishes that

$$\sqrt{m} \left(\frac{1}{m} \sum_{i=1}^m \ln(Z_i) \right) = \frac{W}{\sqrt{m}}$$

is asymptotically normally distributed as $m \rightarrow \infty$. Then since $X = \exp(W)$, the continuous mapping theorem can be used to show that X is asymptotically lognormal. And just as the condition of being identically distributed can sometimes be relaxed for the central limit theorem, the Z_i are not required to be identically distributed for the limiting distribution to be lognormal.

A context in which the lognormal distribution arises as a natural input model is in returns on investments that are compounded. Suppose that $Z_i = (1 + R_i)$, where the R_i are i.i.d. and > -1 . Here Z_i represents the rate of return in the i th period, so the overall rate of return for m periods is $X = \prod_{i=1}^m Z_i$. Thus, the value of an initial investment of \$\$S has value $S \cdot X$ after m periods, which might be modeled as lognormally distributed if m is large.

6.2.1.3 The Weibull and Gamma Distributions

The Weibull and gamma distributions often arise as models of system reliability (time until failure).

There are several equivalent parameterizations of the Weibull distribution, but this is a common one for the pdf and cdf, respectively:

$$f_T(t) = \alpha \beta^{-\alpha} t^{\alpha-1} e^{-(t/\beta)^\alpha} \quad (6.3)$$

$$F_T(t) = 1 - e^{-(t/\beta)^\alpha} \quad (6.4)$$

for $t \geq 0$, where $\alpha > 0$ is called the shape parameter, and $\beta > 0$ is the scale parameter. Notice that when $\alpha = 1$ the Weibull distribution becomes the exponential distribution with mean β .

Let T represent the time until failure of a system. The hazard function of T is a useful way to think about how a distribution represents reliability. The definition of the hazard function is

$$h(t) = \frac{f_T(t)}{1 - F_T(t)} \quad (6.5)$$

which can be interpreted as the system failure rate at time t (numerator), given survival to time t (denominator). For the Weibull,

$$h(t) = \alpha \beta^{-\alpha} t^{\alpha-1}$$

for $t \geq 0$, which is a constant $1/\beta$ if $\alpha = 1$.

By taking the derivative of $h(t)$ with respect to t , we can show that $h(t)$ for the Weibull is increasing for $\alpha > 1$, decreasing for $\alpha < 1$, and constant for $\alpha = 1$. An increasing hazard is appropriate for a system that is more likely to fail as it ages; a decreasing hazard is appropriate for a system that is subject to early failure; while a constant hazard (the exponential special case) means that the failure rate is unaffected by age.

The gamma distribution is similar in many ways to the Weibull; the pdf is

$$f_T(t) = \frac{\beta^{-\alpha} t^{\alpha-1} e^{-(t/\beta)}}{\Gamma(\alpha)} \quad (6.6)$$

for $t \geq 0$, where $\Gamma(\alpha) = \int_0^\infty t^{\alpha-1} e^{-t} dt$ is the gamma function. The parameters α and β have the same interpretation as for the Weibull, with $\alpha = 1$ again giving the

exponential distribution with mean β . Neither the cdf nor the hazard function of the gamma distribution have a simple closed form, but it can be shown that the hazard function may be increasing, decreasing or constant, just as with the Weibull. So how might an understanding of the physics of a process lead to a sensible choice between these two distributions?

Displayed below are the pdfs with scale parameter $\beta = 1$:

Weibull	gamma
$\alpha t^{\alpha-1} e^{-t^\alpha}$	$\Gamma(\alpha)^{-1} t^{\alpha-1} e^{-t}$

The similarities are obvious. But notice that as t increases (t represents the time of failure), the tail of the gamma distribution becomes that of an exponential distribution, no matter what the choice of the shape parameter α . Recall that the hazard function for the exponential is constant. Thus, as the item of interest ages without failing, it eventually has a constant failure rate if it is modeled as a gamma distribution. For the Weibull distribution, however, the shape parameter affects the tail; in particular for $\alpha > 1$ the tail is lighter than exponential, indicating an increasing failure rate even for older items.

Consider the following two examples¹: A new computer operating system is released to the market. Which of the Weibull or gamma distributions might be an appropriate model for the time to the first failure of the operating system? In this case, a probability model would likely be one that has increased risk of a failure initially, but after the software gets “broken in” the risk of a failure decreases to a constant value. Since the risk is constant, one would like to use a gamma distribution, rather than a Weibull distribution, because of the exponential right-tail property. Notice that if we used the Weibull distribution with $\alpha = 1$, then we would have constant failure rate even at product release, which is clearly not appropriate.

As a second example, consider the time to failure of a ball bearing in a machine. In this case, since the ball bearing is physically wearing out, it is not reasonable that as time increases the risk of failure becomes constant. The right-hand tail of a gamma distribution is too heavy and the Weibull distribution is the more appropriate model.

6.2.2 Estimation and Testing

When there is a strong physical basis for choosing a parametric family of distributions, then parameter estimation with a method like maximum likelihood, and goodness-of-fit hypothesis testing of the choice, make sense. These topics are well covered in many textbooks, including Law (2007), and both parameter estimation and testing are included in input-modeling software packages. Therefore, this section provides a few comments to put estimation and testing in context.

¹ These examples were suggested in a personal communication by Dr. Lawrence Leemis of the College of William & Mary.

6.2.2.1 Maximum Likelihood Estimation (MLE)

Suppose that we have a sample of real-world data X_1, X_2, \dots, X_m and we have a strong physical basis for modeling it as being an i.i.d. sample from a distribution with pdf $f^*(x; \theta)$, where θ is a vector of unknown parameters. Then the likelihood function is defined to be

$$L(\theta) = \prod_{i=1}^m f^*(X_i; \theta). \quad (6.7)$$

The MLE is the value of θ that maximizes $L(\theta)$ given the available data:

$$\hat{\theta} = \operatorname{argmax}_{\theta} L(\theta).$$

MLEs have many good statistical properties under the assumption that $f^*(x; \theta)$ is the true, correct distribution. The only reasonable justification for this assumption is that there is strong process physics supporting the selection of $f^*(x; \theta)$.

Notice also that the likelihood function (6.7) is formed by assuming that the data are independent. Even if $f^*(x; \theta)$ is an appropriate parametric family of distributions for X , (6.7) is not the likelihood function if the data are dependent (in which case the likelihood function involves the full joint distribution). Preliminary examination of the data for signs of dependence is always worthwhile.

6.2.2.2 Goodness-of-Fit Testing

The premise behind goodness-of-fit testing is that there is a true, correct distribution f^* with true, correct parameter values θ^* . The hypothesis test is then

$$\begin{aligned} H_0: f(x; \hat{\theta}) &= f^*(x; \theta^*) \\ H_1: f(x; \hat{\theta}) &\neq f^*(x; \theta^*). \end{aligned}$$

Because real-world data do not come from probability distributions, we know before executing the test that the null hypothesis is false. However, if we have a strong basis for choosing $f(x; \theta)$ and have used an appropriate parameter estimator $\hat{\theta}$, then such a test is useful because it may alert us to egregious departures of the data from the proposed input model. A few additional points are worth keeping in mind:

- When the sample size is small, goodness-of-fit tests often have low power and thus may accept (formally, fail to reject) many possible parametric families of distributions. As discussed above, the Weibull and gamma families are very similar, and therefore may be indistinguishable when m is small. Thus, accepting the null hypothesis should not be interpreted as proof that the correct, or best, choice has been made.
- When the sample size is large, goodness-of-fit tests will often reject every choice of parametric distribution family because real-world data do not come from

probability distributions, and when m is large, then the test has enough power to recognize this. Thus, rejecting the null hypothesis does not necessarily mean that the chosen distribution should be discarded.

- Rather than report an accept/reject decision, most input modeling software provides either the value of the test statistic itself (a small value is better for most tests), or a p -value. The p -value is the Type I error level at which one would just reject the null hypothesis. Thus, a large p -value (say >0.1) suggests accepting the chosen distribution rather than taking a large risk by rejecting it.
- There are many goodness-of-fit tests, each sensitive to different sorts of departures from the null hypothesis. Thus, it is entirely possible that one test may give a large p -value (suggesting that you accept the chosen distribution), and another may give a small p -value (suggesting you reject the chosen distribution) for the same data. In fact, the very popular chi-squared test may give both conclusions depending upon how many bins into which the data are divided. *For this reason, automated selection of an input model by fitting many of them and choosing the one with the smallest value of a particular test statistic or the largest p -value for that test seems questionable.* If there is no strong physical support for a particular distribution, then the matching paradigm (discussed below) may be more relevant.

Goodness-of-fit tests represent fit or lack of fit by a summary statistic. Distribution fitting software usually also contains graphical tools for visually evaluating fit, and these give a more comprehensive picture. The most popular is the density-histogram plot, but your perception of fit for that plot can be influenced by how the data are grouped to form the histogram. A better tool that does not require grouping is the quantile-quantile (Q-Q) plot. For a random variable with continuous, increasing cdf F_X its q -quantile (for $0 < q < 1$) is $F_X^{-1}(q)$, the inverse cdf evaluated at q . Recall that we used the inverse cdf for variate generation, $X = F_X^{-1}(U)$ with $U \sim U(0,1)$. The Q-Q plot plots the sorted input data $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(m)}$ against the fitted inverse cdf \hat{F}^{-1} at what might be called a “perfectly generated sample”

$$\hat{F}^{-1}\left(\frac{0.5}{m}\right) < \hat{F}^{-1}\left(\frac{1.5}{m}\right) < \dots < \hat{F}^{-1}\left(\frac{m-0.5}{m}\right).$$

Stated more succinctly, it plots $X_{(i)}$ vs. $\hat{F}^{-1}((i-0.5)/m)$, and a good fit is indicated by $X_{(i)} \approx \hat{F}^{-1}((i-0.5)/m)$, $i = 1, 2, \dots, m$. Deviation from an approximate 45° line indicates not only lack of fit, but also where it occurs.

To illustrate how it works, Fig. 6.1 shows Q-Q plots for a set of data that actually are normally distributed, but for which both uniform and normal distributions have been fit. The uniform distribution is far too heavy in its tails, which the plot reveals in the way that the tails deviate more severely from a line.

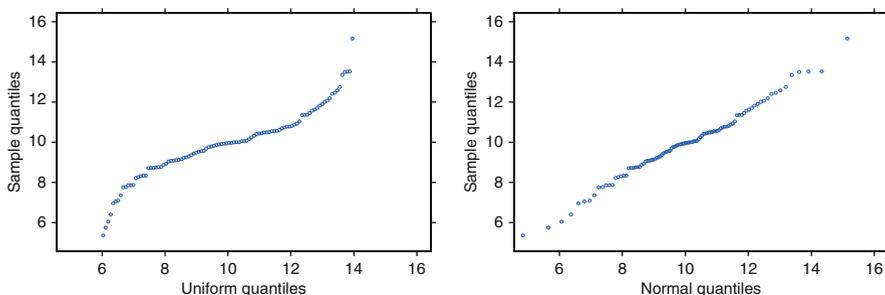


Fig. 6.1 Q-Q plot resulting from fitting a uniform (*left*) and normal (*right*) distribution to a set of input data

6.2.2.3 Known and Unknown Bounds

Many standard distributions have their support on a fixed domain; for instance, the lognormal, Weibull and gamma distributions are defined on $[0, \infty)$, while the beta distribution is defined on $[0, 1]$.

A random variable X' describing outcomes on $[0, \infty)$ can be shifted to $[a, \infty)$ simply by adding a constant: $X = a + X'$. This changes the mean but not the variance. A random variable Y' describing outcomes on $[0, 1]$ can be shifted and scaled to $[a, b]$ by the transformation $Y = a + (b - a)Y'$; this changes both the mean and variance.

Obviously these transformations can be inverted to convert a random variable to a standard domain: $X' = X - a$ and $Y' = (Y - a)/(b - a)$. This is important if we encounter a situation in which there are known bounds. For instance, we might want to model a processing time using a distribution with standard domain $[0, \infty)$, but we know that the process cannot take less than 4 min. To fit the observed data X'_1, X'_2, \dots, X'_m , we should first transform it by subtracting 4 (e.g., $X_i = X'_i - 4$), fit the distribution to the transformed data, and then add 4 to each random variate generated.

Using known bounds is not the same as trying to infer an unknown bound. Distribution fitting software will often treat each distribution on $[0, \infty)$ as having an additional lower-bound parameter a , and each distribution on $[0, 1]$ as having two additional upper and lower-bound parameters a, b ; these additional parameters are then included in the fitting. Sometimes the bounds are estimated heuristically (e.g., \hat{a} is the smallest observation), and sometimes they are included in the likelihood function. *Known bounds should typically be used instead of estimated bounds, if available.* A known bound brings in additional information that may not be present in the data, and estimated bounds may substantially over (under) estimate lower (upper) bounds, particularly in small samples.

6.2.3 Matching Properties of Univariate Distributions

When there is not a strong physical basis for choosing a particular family of distributions, then an alternative to trying to infer the “true” distribution is to use a very flexible family of distributions and obtain values for its parameters so that it closely matches properties of the data. In this section we describe two common matching methods and present one flexible family.

Recall that for the $M/G/1$ queue the steady-state expected waiting time depends only on the mean and variance of the service-time distribution; thus, any nonnegative distribution that matches the correct mean and variance will give the right simulation results. Typically more than just the mean and variance matter, which is why there are many two-parameter distributions (Weibull, gamma, lognormal, etc.) from which to choose. Fortunately, it is often the case that if we can match the first four moments (defined below) of the input random variable correctly, then the particular family of distributions is not as important.²

Let X be a random variable with distribution F_X . The k th moment of X is

$$E(X^k), \quad k = 1, 2, \dots$$

Of more use for input modeling are the standardized central moments, specifically the first four:

$$\begin{aligned} \mu_X &= E(X) && \text{mean} \\ \sigma_X^2 &= E[(X - \mu_X)^2] && \text{variance} \\ \alpha_3 &= E[(X - \mu_X)^3] / \sigma_X^3 && \text{skewness} \\ \alpha_4 &= E[(X - \mu_X)^4] / \sigma_X^4 && \text{kurtosis} \end{aligned} \quad (6.8)$$

The skewness is a measure of symmetry (symmetric distributions have $\alpha_3 = 0$), while the kurtosis is a measure of tail weight (the normal distribution has kurtosis $\alpha_4 = 3$; sometimes $\alpha_4 - 3$ is called the excess kurtosis). If α_3 and α_4 are finite, then $\alpha_4 > 1 + \alpha_3^2$ and the inequality is strict. This relationship defines a plane of feasible (α_3^2, α_4) values.

Suppose that the first four central moments of X , $(\mu_X, \sigma_X^2, \alpha_3, \alpha_4)$, exist. Let μ and σ^2 be a desired mean and variance, and define a new random variable

$$X' = \mu + \sigma \left(\frac{X - \mu_X}{\sigma_X} \right). \quad (6.9)$$

Then it is easy to show that X' has mean μ and variance σ^2 ; in the exercises you are asked to show that X and X' have the same skewness and kurtosis. Thus, any random variable X can be transformed to match a desired mean and variance without altering its skewness and kurtosis. In matching, α_3 and α_4 are the challenges.

For a parametric distribution $F(\cdot; \theta)$, its central moments are functions of θ ; denote these as $\mu(\theta)$, $\sigma^2(\theta)$, $\alpha_3(\theta)$ and $\alpha_4(\theta)$, respectively. We call a parametric distribution “flexible” if its parameters θ allow coverage of a substantial portion

² The most common exceptions are when system performance depends critically on the behavior of the extreme tail of the distribution, or when one or more of the first four moments are infinite.

of the feasible (α_3^2, α_4) plane. As (6.9) shows, the mean and variance of F are not critical since we can scale and shift any distribution to have the desired first two central moments.

Many well-known distributions provide little flexibility; for instance the normal distribution has $\alpha_3 = 0$ and $\alpha_4 = 3$; the exponential distribution has $\alpha_3 = 2$ and $\alpha_4 = 9$; while the gamma distribution (6.6) has a little flexibility with $\alpha_3 = 2/\sqrt{\alpha}$ and $\alpha_4 = 3 + 6/\alpha$. We describe some more flexible distributions below.

The sample standardized central moments of the data are

$$\begin{aligned}\bar{X} &= \frac{1}{m} \sum_{i=1}^m X_i \\ \hat{\sigma}^2 &= \frac{1}{m} \sum_{i=1}^m (X_i - \bar{X})^2 \\ \hat{\alpha}_3 &= \frac{1}{m} \sum_{i=1}^m (X_i - \bar{X})^3 / \hat{\sigma}^3 \\ \hat{\alpha}_4 &= \frac{1}{m} \sum_{i=1}^m (X_i - \bar{X})^4 / \hat{\sigma}^4.\end{aligned}$$

To match the moments, we solve the system of equations

$$\begin{aligned}\mu(\theta) &= \bar{X} \\ \sigma^2(\theta) &= \hat{\sigma}^2 \\ \alpha_3(\theta) &= \hat{\alpha}_3 \\ \alpha_4(\theta) &= \hat{\alpha}_4\end{aligned}$$

for θ ; denote the solution by $\hat{\theta}_M$. Unless we bring in some other type of information, we need at least four parameters to match four moments.

Provided that the parametric family is flexible enough to match the moments exactly, then using $F(\cdot; \hat{\theta}_M)$ as the input model guarantees that the simulation input has the same central moments as the observed data X_1, X_2, \dots, X_m ; that is, they match. For this reason the “method of moments” is an appealing approach that is widely used. Notice that one need not always fit four moments. For instance, the gamma distribution has $\theta = (\alpha, \beta)$ with $\mu(\theta) = \alpha\beta$ and $\sigma^2(\theta) = \alpha\beta^2$, so it can be fit by solving $\hat{\alpha}\hat{\beta} = \bar{X}$ and $\hat{\alpha}\hat{\beta}^2 = \hat{\sigma}^2$.

A drawback of the method of moments is that there is no guarantee that the fitted distribution $F(\cdot; \hat{\theta}_M)$ resembles the empirical distribution of the data. The next approach tries to achieve a closer match between these two.

A second matching approach exploits the following fact: If X is a random variable with strictly increasing cdf F_X , then $U = F_X(X)$ has a $U(0, 1)$ distribution. Exercise 36 asks you to show that this is indeed the case, but it seems reasonable since it reverses the inverse cdf method of variate generation. This fact implies that if X_1, X_2, \dots, X_m are i.i.d. with continuous distribution F_X , then $U_i = F_X(X_i), i =$

$1, 2, \dots, m$ are i.i.d. $U(0, 1)$. Going one step further, let $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(m)}$ be the order statistics of the sample of data, and let $U_{(i)} = F_X(X_{(i)})$, $i = 1, 2, \dots, m$. Then $U_{(i)}$ has the distribution of the i th order statistic of m i.i.d. $U(0, 1)$ random variables, for which it is known that

$$E(U_{(i)}) = \frac{i}{m+1}.$$

If we want to match the behavior of X_1, X_2, \dots, X_m with a parametric distribution $F(\cdot; \theta)$, then another approach that is different from the method of moments is to set

$$\hat{\theta}_U = \operatorname{argmin}_{\theta} \sum_{i=1}^m w(i) \left(F(X_{(i)}; \theta) - \frac{i}{m+1} \right)^2 \quad (6.10)$$

where $w(i)$ is a positive weight. The weights can be chosen based on a number of criteria; common choices are $w(i) = 1$ and $w(i) = (m+2)(m+1)^2/[i(m-i+1)]$. The latter weight is $1/\operatorname{Var}(U_{(i)})$ so that it gives order statistics with smaller variance more weight.

What does this fit achieve? Notice that if the fit was perfect, then

$$F(X_{(i)}; \hat{\theta}_U) = \frac{i}{m+1}.$$

Compare this with the empirical cumulative probability of the real-world data:

$$\frac{\#\{X_j \leq X_{(i)}\}}{m} = \frac{i}{m}.$$

Thus, with a perfect fit the cumulative probabilities of the fitted distribution at the data points $F(X_i; \hat{\theta}_U)$, $i = 1, 2, \dots, m$ are nearly the same as the cumulative probabilities of the data themselves. Therefore, F should tend to look like the empirical distribution of the data. Exercise 5 asks you to show why it is better to fit to $i/(m+1)$ rather than i/m .

A similar fit is obtained by using the inverse cdf. Let $Q(u; \theta) = F^{-1}(u; \theta)$, the inverse cdf. If $E[Q(U_{(i)}; \theta)]$ is easy to evaluate, then an alternative to (6.10) is to fit directly to the order statistics of X :

$$\hat{\theta}_O = \operatorname{argmin}_{\theta} \sum_{i=1}^m (X_{(i)} - E[Q(U_{(i)}; \theta)])^2. \quad (6.11)$$

We refer to both $\hat{\theta}_U$ and $\hat{\theta}_O$ as *least-squares fits*.

To employ either the method of moments or least squares matching, we need flexible families of distributions. The flexible distributions that have seen significant application in simulation are Johnson's translation system (see, for instance Swain et al. 1988) and the generalized lambda distribution (GLD; see for instance Karian and Dudewicz 2000). Johnson's translation system is comprised of four transformations of the normal distribution and it covers the entire feasible (α_3^2, α_4) plane; the

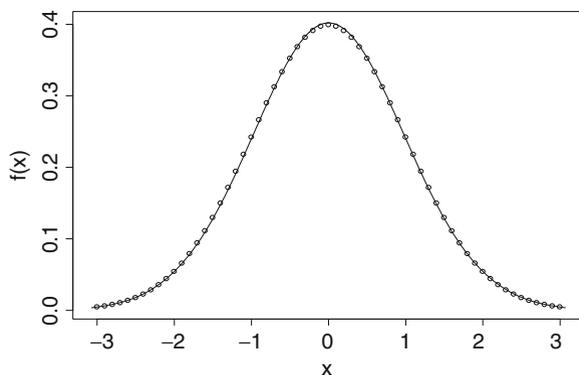


Fig. 6.2 GLD approximation (solid line) to points on the standard normal pdf

GLD is a single function that covers a large portion of it. We present the GLD here because it is easy to describe as well as being useful.

The GLD has four parameters $\theta = (\lambda_1, \lambda_2, \lambda_3, \lambda_4)$, and is most easily represented via its inverse cdf

$$Q(u; \theta) = \lambda_1 + \frac{u^{\lambda_3} + (1-u)^{\lambda_4}}{\lambda_2} \quad (6.12)$$

for $0 \leq u \leq 1$. Therefore, variate generation is immediate. The parameter λ_1 controls location, λ_2 controls scale, and skewness and kurtosis are determined jointly by λ_3 and λ_4 .

There is no physical basis for choosing the GLD; its usefulness comes from its flexibility and the ease of random-variate generation. Unfortunately fitting is not as easy, and this is typical of flexible distributions with four or more parameters. Fitting via the method of moments (Karian and Dudewicz 2000) and least squares (using formulation (6.11), see Lakhany and Mausser (2000) and references therein) are possible, but both involve a numerical search for parameters since there are no closed-form solutions. In the Appendix we give expressions for the necessary moments and expectations.

Figure 6.2 shows the GLD approximation for the $N(0, 1)$ distribution, where the circles are points on the normal pdf, while the solid line is the GLD approximation. The standard normal has moments $\mu_X = 0$, $\sigma_X^2 = 1$, $\alpha_3 = 0$ and $\alpha_4 = 3$. Clearly out to three standard deviations from the mean the approximation is quite good. Figure 6.3 shows the GLD approximation to the exponential distribution with mean 1, which has moments $\mu_X = 1$, $\sigma_X^2 = 1$, $\alpha_3 = 2$ and $\alpha_4 = 9$. While it has the correct moments, the GLD approximation has a different shape as x approaches 0, and in fact assigns positive (but very, very small) probability to negative x . Matching moments does not guarantee a particular distribution shape.

The GLD parameters $(\lambda_1, \lambda_2, \lambda_3, \lambda_4)$ for these examples were obtained from Karian and Dudewicz (2000); for the normal approximation they are $(0, 0.1975, 0.1349, 0.1349)$; while for the exponential approximation they are $(0.006862, -0.0010805, -0.4072 \times 10^{-5}, -0.001076)$.

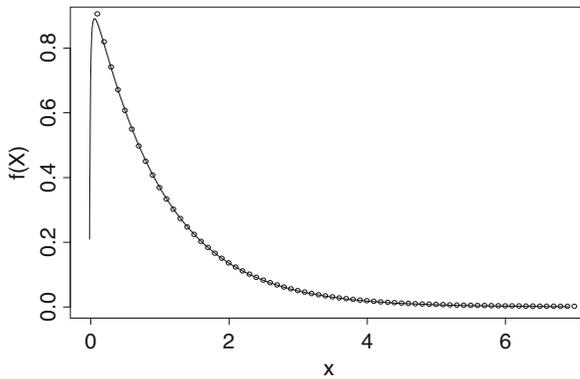


Fig. 6.3 GLD approximation (solid line) to points on the exponential pdf with mean 1

6.2.4 Empirical Distributions

As in the previous sections, we have a sample of real-world data X_1, X_2, \dots, X_m that we believe to be i.i.d. observations from some unknown distribution F_X . When there is no strong physical basis for choosing a particular parametric family of distributions, and no parametric distribution provides an adequate fit to the data, then it makes sense to consider using the data themselves. The disadvantages to using the data directly will be apparent shortly.

The empirical cdf (ecdf) is a nonparametric input model defined as

$$\hat{F}(x) = \frac{1}{m} \sum_{i=1}^m I(X_i \leq x) \tag{6.13}$$

for all $-\infty < x < \infty$, where $I(\cdot)$ is the indicator function. Let $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(m)}$ be the sorted values, and let \hat{X} be a random variable with distribution \hat{F} . The ecdf places probability mass $1/m$ on each observed value:

$$\begin{aligned} \Pr \left\{ \hat{X} = X_{(i)} \mid X_1, \dots, X_m \right\} &= \Pr \left\{ \hat{X} \leq X_{(i)} \mid X_1, \dots, X_m \right\} - \Pr \left\{ \hat{X} < X_{(i)} \mid X_1, \dots, X_m \right\} \\ &= \frac{i}{m} - \frac{(i-1)}{m} = \frac{1}{m}. \end{aligned}$$

This makes variate generation easy via the inverse cdf:

1. Generate $U \sim U(0, 1)$
2. Set $i = \lceil mU \rceil$
3. Return $\hat{X} = X_{(i)}$

The ecdf has several appealing properties. First, as an estimator of the true distribution F_X it is unbiased:

$$\mathbb{E}(\widehat{F}(x)) = \mathbb{E}\left(\frac{1}{m} \sum_{i=1}^m I(X_i \leq x)\right) = \frac{m}{m} \mathbb{E}(I(X_1 \leq x)) = F_X(x).$$

Notice that the expectation is with respect to all possible real-world samples from F_X . Further, using the strong law of large numbers it is easy to show that $\widehat{F}(x) \xrightarrow{a.s.} F_X(x)$ pointwise as $m \rightarrow \infty$, since $\widehat{F}(x)$ is the average of i.i.d. observations $I(X_i \leq x)$, $i = 1, 2, \dots, m$.

A second feature of the ecdf is that if we use random variates \widehat{X} in our simulation, then they will have properties that match the sample properties of the observed real-world data; for instance,

$$\mathbb{E}(\widehat{X} \mid X_1, \dots, X_m) = \sum_{i=1}^m X_{(i)} \frac{1}{m} = \bar{X}.$$

Here the expectation is with respect to the ecdf \widehat{F} , and it shows that the expected value of a random variable with distribution \widehat{F} is the sample mean of the data that created it.

These appealing properties of the ecdf contrast with two undesirable ones: only the observed values X_1, X_2, \dots, X_m have positive probability, and only values between $[X_{(1)}, X_{(m)}]$ will be realized. In other words, it is a finite, discrete distribution. This is particularly troubling if X is known to be continuous valued, or it is expected to have a long tail of extreme but low probability values.

A standard way to obtain a continuous distribution is to linearly interpolate the ecdf between the data points:

$$\tilde{F}(x) = \begin{cases} 0, & x < X_{(1)} \\ \frac{i-1}{m-1} + \frac{x - X_{(i)}}{(m-1)(X_{(i+1)} - X_{(i)})}, & X_{(i)} \leq x < X_{(i+1)} \\ 1, & x \geq X_{(m)}. \end{cases} \quad (6.14)$$

The linearly interpolated ecdf fills in between the observed data; however, the support is still restricted to $[X_{(1)}, X_{(m)}]$. Bratley et al. (1987) describe a linearly interpolated ecdf with an exponential tail added.

Because the linearly interpolated ecdf is piecewise linear, variate generation via inversion is easy:

1. Generate $U \sim U(0, 1)$
2. Set $i = \lceil (m-1)U \rceil$
3. Return $\tilde{X} = X_{(i)} + (m-1)(X_{(i+1)} - X_{(i)}) \left(U - \frac{i-1}{m-1} \right)$

Unfortunately, \tilde{F} is biased for F_X , although $\tilde{F}(x) \xrightarrow{a.s.} F_X(x)$ as $m \rightarrow \infty$. The proof of this is instructive because it shows that various smoothing schemes can be used without losing asymptotic consistency.

Proof. For any fixed x , let

$$\tilde{F}(x) = \frac{1}{m-1} \sum_{i=1}^m I(X_i \leq x) = \frac{m}{m-1} \hat{F}(x).$$

Then clearly $\tilde{F}(x) \xrightarrow{a.s.} F_X(x)$ because $m/(m-1) \rightarrow 1$. Notice also that for $X_{(i)} \leq x < X_{(i+1)}$ we have

$$0 \leq \frac{x - X_{(i)}}{(m-1)(X_{(i+1)} - X_{(i)})} < \frac{1}{m-1}.$$

Therefore,

$$\tilde{F}(x) - \frac{1}{m-1} \leq \tilde{F}(x) < \tilde{F}(x).$$

Since both the lower and upper bounds on $\tilde{F}(x)$ converge a.s. to $F_X(x)$, so does $\tilde{F}(x)$. \square

Next consider the random variable \tilde{X} . Exercise 10 asks you to show that $E(\tilde{X} | X_1, \dots, X_m) \neq \bar{X}$, but the discrepancy decreases with m . Thus, the linearly interpolated ecdf does not perfectly match the sample properties of the data when m is small. Nevertheless, smoothing is typically worthwhile when the input is expected to be continuous valued.

6.2.5 Direct Use of Input Data

Arrivals to the fax center described in Sect. 4.6 were modeled by a nonstationary Poisson arrival process fit to actual arrival data. Fax arrival data are relatively easy to collect since each incoming fax has a time stamp printed on it by the fax machine, and after the fax has been entered we know if it was classified as “simple” or not. Days, months or even years of such data might be available. This suggests that we could avoid input modeling by driving the simulation with the actual observed arrival data. Functionally, the simulation would read the arrival times and fax types from a file of actual data rather than generating them from an input model. Notice that this is different from using the empirical distribution, which resamples the data rather than using it exactly as it was observed. This distinction is important to understanding the advantages and disadvantages of direct data use:

Advantages: Direct use might capture features missed by an input model, even the empirical distribution. For instance, if our input model treats a process as i.i.d., but there is actually some difficult-to-observe dependence or nonstationary behavior, then this will be reflected in the data but not the input model. In

fact, probabilistic input models will seldom be rich enough to capture all of the complexity of a real process.

Disadvantages: Clearly the simulation run length is limited by the amount of data we have; increasing the run length or number of replications is not possible. And the statistical statements we can make are subtle. For instance, in the fax center simulation we could make statements such as “had we staffed at this level during the period when we recorded fax arrivals, our performance would have been. . . .” Statements about long-run performance, however, are more problematic. Consider a system that is affected by rare occurrences in the input (e.g., excessively large insurance claims). However, the system can tolerate isolated occurrences of the rare input value, but not two or more in quick succession. If multiple rare occurrences in close proximity never happened in the observed data, then of course they will never happen in a simulation that directly uses the data; the simulation estimate of the chance of disaster will be 0. However, if we use a probabilistic input model (e.g., modeling insurance claims as being i.i.d. with a fitted distribution F), then over a long enough simulation we will observe multiple rare input values together and recognize the vulnerability of the system.

There is no easy rule for deciding when direct use of the data is appropriate and when fitting an input model is better. When we have large quantities of data that we believe are representative, and we have no reason to believe that the system is dramatically affected by unobserved behavior in the input, then direct data use is reasonable. When we want to make broader statements that go beyond just the input data we saw, then input modeling is essential.

6.2.6 *Input Modeling Without Data*

Many simulations are undertaken without having relevant input data available. When this happens we try to exploit knowledge about the process to produce reasonable input models, and sensitivity analysis to assess how much they matter.

As discussed in Sect. 6.2, the physical nature of the input process may align well with certain parametric distribution families. This sort of thinking is particularly helpful when there are no data. The difficulty then becomes assigning values to the distribution’s parameters. One approach is to obtain subjective estimates from people familiar with the process. However, care must be taken to ask for information in a way that is likely to elicit useful responses. Here is an example.

Suppose the nature of the process suggests use of a normal distribution. Therefore, numerical values for the mean μ and variance σ^2 are required. People are good at providing “typical” values, although a clear distinction between the mean (average) value and the most likely value is not common. Fortunately, for the normal distribution they are the same thing. The variance, however, is very difficult to estimate based on experience alone (without having looked at data). Here are three ways that a value for the variance might be obtained:

Estimate the standard deviation: Unlike the variance, the standard deviation σ is measured in the natural units of the input process, the same units as the mean. Someone familiar with a process might be able to provide the average deviation around the mean, which is an intuitive way to ask for the standard deviation.

Estimate an extreme deviation: If someone familiar with the process can provide an extreme, but still feasible, deviation from the mean, then this could be interpreted as, say, 3σ . Sensitivity analysis could assume multiples of σ that are greater or less than 3.

Relate to a known process: If there is a similar process on which data are available, then one can ask if the other process is more or less variable, and by what percentage. If the process with data has standard deviation S , and the process without data is thought to be 10% less variable, then $\sigma = 0.9S$. Sensitivity analysis could assume values greater or less than 0.9.

The key idea is to translate the parameters of the distribution into terms that are understandable by those familiar with the process.

When there is no strong physical basis for selecting a particular family of distributions, then a distribution like the triangular may be useful. The triangular distribution was not created with a specific physical process in mind; instead it is designed to match three easy-to-specify characteristics: minimum value (a), most likely value (b) and maximum value (c). The pdf of the triangular distribution is

$$f(x) = \begin{cases} \frac{2(x-a)}{(b-a)(c-a)}, & a \leq x \leq b \\ \frac{2(c-x)}{(c-b)(c-a)}, & b \leq x \leq c \\ 0, & \text{elsewhere.} \end{cases} \quad (6.15)$$

Notice that triangular distributions with $a = b$ or $b = c$ are also legitimate distributions where one of the extremes is also the most likely value.

The triangular distribution is typically superior to the uniform distribution on the minimum and maximum values, because there are few real processes for which the extremes (a and c) are just as likely as more central values. One justification for choosing the uniform, however, is that it represents maximum uncertainty on values between a and c .

The uniform and triangular distributions are not the only ones that can be fit from this sort of information. For instance, the beta distribution can also be parameterized by the minimum, maximum and most likely values. However, there is not a unique beta distribution to represent these characteristics, while the triangular is unambiguously defined. Once there is no longer a physical basis for the distribution choice, then the triangular distribution is as good as any other choice, and it is easy to vary the extremes and the most likely value to check sensitivity.

6.3 Multivariate Input Processes

The previous section described input processes consisting of independent and identically distributed observations; we referred to these as univariate input processes because specification of the marginal distribution is all that is required. This section covers two types of multivariate input processes that are frequently needed in practice: nonstationary arrival processes and random vectors.

Examples of (possibly nonstationary) arrival processes include the arrival of customers to a shopping mall; the arrivals of e-mails to a mail server; and the arrivals of claims to an insurance company. Although we can think of defining the joint distribution of all of the arrival times, it is convenient for discrete-event simulation to describe the *times between arrivals*; this permits the simulation to advance from one arrival to the next by simply scheduling the next arrival.

Examples of random-vector input processes include multiple characteristics of the same customer (e.g., gender, age, income and occupation); annual sales of related products (e.g., new cars, new car tires, custom floor mats and GPS devices); and returns on various financial assets (e.g., values of short-term bonds, long-term bonds and stocks). An input model for a random vector must specify the marginal distribution of each component of the vector and the dependence among them.

There is a vast literature on nonstationary arrival processes that are Poisson, and on multivariate probability distributions that are useful as input models. Good starting points are Leemis (2006) for Poisson arrival processes, and Johnson (1987) and Biller and Ghosh (2006) for random vectors.

Just as was the case with univariate input models, the physical basis of the real process can provide a justification for the choice of multivariate input model. However, models that are obtained in this way tend to be somewhat restrictive. Many well-known multivariate probability distributions have all of their component marginal distributions from the same family. For instance, all marginals of the multivariate normal distribution are normal, making it unsuitable for a situation in which one marginal should be continuous valued, while another is discrete.

For this reason transformation-based approaches for constructing multivariate inputs processes have proven useful in simulation. By “transformation-based approach” we mean that the input process we want is obtained by transforming a more basic input process. To be useful, the base process should have easily controllable characteristics, which the transformation preserves, while the transformation itself allows matching some other desired features of the input process. To be specific, we will obtain nonstationary arrival processes by transforming i.i.d. interarrival times to obtain the desired arrival rate while preserving a measure of variability; and we will obtain random vectors by transforming a vector of dependent $U(0, 1)$ random variables to obtain the desired marginal distributions.

The transformation-based approaches we describe provide a general framework for creating arrival processes with desired arrival rate and variability, and random vectors with given marginals and correlations. There is no claim that the physical basis of the real process justifies the transformation. Stated differently, the transformation-based approaches are tools for *matching* characteristics.

A convenient feature of this approach is that random-variate generation follows directly: generate the base process, then transform it to obtain inputs. For this reason we will present both input modeling and random-variate generation together in this section.

Time series are a third class of multivariate input process, but they are not covered in this book; a brief introduction with references can be found in Biller and Ghosh (2006). A time-series input process represents a sequence of (possibly vector) random variables that are dependent in sequence, for example the sequence of weekly order quantities that a grocery store chain places for a sports drink.

6.3.1 Nonstationary Arrival Processes

Arrival processes are often primary inputs to computer simulation of real-world systems. The parking lot in Sect. 3.1 had an arrival process representing cars, while the arrivals to the hospital reception in Sect. 3.2 were patients and visitors. A characteristic of these two examples, and many real arrival processes, is that the arrivals are not organized or predictable. In such cases, the usual input modeling paradigm is to characterize the distribution of the time between arrivals, also called the *interarrival time*. The interarrival-time approach does not apply to scheduled arrivals, such as patients to a doctor's office, and may or may not be a good model for arrivals targeting a particular event time, such as fans arriving to a basketball game; in both of these cases the total number of arrivals is fixed (number of appointments or number of tickets sold, respectively) and the arrivals are supposed to occur at or around a specific time. See Exercises 11 and 22.

Renewal processes are the simplest form of random arrivals: In a renewal arrival process, $\tilde{A}_1, \tilde{A}_2, \dots$ are i.i.d. nonnegative random variables with distribution G that represent the interarrival times of successive entities. The arrivals to the hospital reception were modeled in this way. In addition to the interarrival times, two related stochastic processes are

$$\tilde{S}_n = \begin{cases} 0, & n = 0 \\ \sum_{i=1}^n \tilde{A}_i, & n = 1, 2, \dots \end{cases}$$

$$\tilde{N}(t) = \max\{n \geq 0 : \tilde{S}_n \leq t\}.$$

In words, \tilde{S}_n is the time of the n th arrival, while $\tilde{N}(t)$ is the number of arrivals by time $t \geq 0$. The continuous-time process \tilde{N} is also called the *arrival-counting process*. Renewal arrival processes are easy to simulate, since an arrival at time \tilde{S}_n triggers scheduling the next arrival to occur at time $\tilde{S}_{n+1} = \tilde{S}_n + \tilde{A}_{n+1}$. If the interarrival times have finite mean and variance, and we let $\tilde{\lambda} = 1/E(\tilde{A}_i)$, then

$$\frac{E(\tilde{N}(t))}{t} \longrightarrow \tilde{\lambda}$$

as $t \rightarrow \infty$, which is why $\tilde{\lambda}$ is interpreted as the arrival rate of the process (see, e.g., Kulkarni 1995). From here on we will assume that the interarrival times \tilde{A}_i are continuous-valued and have a density function, which is appropriate for arrivals that occur in continuous time.³

The parking lot did not have a renewal arrival process; instead it had an arrival process with a time-varying arrival rate $\lambda(t) = 1,000 + 100 \sin(\pi t/12)$. Such arrival processes are called *nonstationary*, and are the focus of this section. We will, however, simulate nonstationary arrival processes by transforming renewal arrival processes in one of two ways: either by stretching (to space arrivals farther apart) or compressing (to force arrivals closer together) the time between renewal arrivals, or by selectively ignoring (“thinning”) some renewal arrivals. To be concrete we use the example of the arrival of e-mails to a mail server.

To do this we need the concept of an *equilibrium renewal process*. The only difference between an equilibrium renewal process and the renewal process defined above is that the first interarrival time \tilde{A}_1 has distribution G_e , where

$$G_e(t) = \Pr\{\tilde{A}_1 \leq t\} = \tilde{\lambda} \int_0^t (1 - G(s)) ds. \quad (6.16)$$

Think of this as the distribution of the time until the next arrival if we started observing the renewal arrival process at an arbitrary point in time.⁴ If we initialize the renewal process in this way, then it can be shown that

$$\frac{E(\tilde{N}(t))}{t} = \tilde{\lambda} \quad (6.17)$$

for all $t \geq 0$, not just in the limit (Kulkarni 1995).

Now consider arrival processes with time-varying arrival rates. Let $N(t)$ be the arrival-counting process of a (potentially) nonstationary arrival process. To make the concept of “arrival rate” precise, define

$$\Lambda(t) = E(N(t))$$

the expected number of arrivals by time t , and the arrival rate by

$$\lambda(t) = \frac{d}{dt} \Lambda(t)$$

when Λ is differentiable. The function $\Lambda(t)$ is also called the *integrated rate function*. For an equilibrium renewal process, Eq. (6.17) implies that $\Lambda(t) = \tilde{\lambda}t$ and $\lambda(t) = \tilde{\lambda}$; to simulate nonstationary arrival processes we allow more general $\Lambda(t)$ and $\lambda(t)$.

³ It is certainly possible to have discrete-time arrival processes when activities are synchronized to a clock, as in some computer networks.

⁴ If G is an exponential distribution, then $G_e = G$ due to the memoryless property of the exponential distribution. Exercise 25 asks you to prove this.

6.3.1.1 Inverting $\Lambda(t)$

Suppose that the time-varying behavior that we want is specified through $\Lambda(t)$ (which might be obtained by integrating $\lambda(t)$). Let \tilde{S}_n be an equilibrium renewal arrival process with arrival rate $\tilde{\lambda} = 1$. Define a nonstationary arrival process with arrival times S_n , interarrival times A_n , and arrival-counting process $N(t)$ as follows:

1. Set index $n = 1$ and $\tilde{S}_0 = 0$
2. Generate \tilde{A}_n
3. Let
 - a. $\tilde{S}_n = \tilde{S}_{n-1} + \tilde{A}_n$
 - b. $S_n = \Lambda^{-1}(\tilde{S}_n)$
 - c. $A_n = S_n - S_{n-1}$
4. $n = n + 1$
5. Go to Step 2

Let $N(t) = \max\{n \geq 0 : S_n \leq t\}$. Notice that $\Lambda(t)$ provides a change of time scale: If it is time s for \tilde{N} then it is time $\Lambda^{-1}(s)$ for N ; similarly, if it is time t for N , then it is time $\Lambda(t)$ for \tilde{N} . Therefore, $N(t) = \tilde{N}(\Lambda(t))$ and

$$\begin{aligned} E(N(t)) &= E[E(N(t) | \tilde{N}(\Lambda(t)))] \\ &= E[\tilde{N}(\Lambda(t))] \\ &= 1 \cdot \Lambda(t) = \Lambda(t) \end{aligned}$$

remembering that \tilde{N} is an equilibrium renewal process with rate 1. Thus, Λ^{-1} transforms the arrival times of an equilibrium renewal process with arrival rate 1 into a nonstationary arrival process with rate $\lambda(t) = d\Lambda(t)/dt$. Inversion provides a very simple way to obtain a nonstationary process with given arrival rate, provided Λ is easily invertible.

Figure 6.4 illustrates the inversion method for an arrival process with $\lambda(t) = 2t$, so that $\Lambda(t) = t^2$. Since the rate is increasing in time, we expect arrivals to be more and more densely packed as t increases. In the figure, the vertical axis represents the arrival times in a rate-1 base process, which is still a random process but with a mean time between arrivals of 1. These times are mapped into arrival times in the nonstationary process by $\Lambda^{-1}(s) = \sqrt{s}$, as shown on the horizontal axis. Notice that as t increases the arrival rate is clearly increasing as well.

6.3.1.2 Thinning for $\lambda(t)$

Suppose now that the time-varying behavior that we want is specified through the arrival rate $\lambda(t)$ (which might be obtained by differentiating $\Lambda(t)$). Let \tilde{S}_n be an equilibrium renewal arrival process with arrival rate $\tilde{\lambda} = \max_t \lambda(t)$, which we assume is finite. The idea behind thinning is to generate a stationary process of potential arrivals at the maximum rate, then randomly delete or “thin” some of the

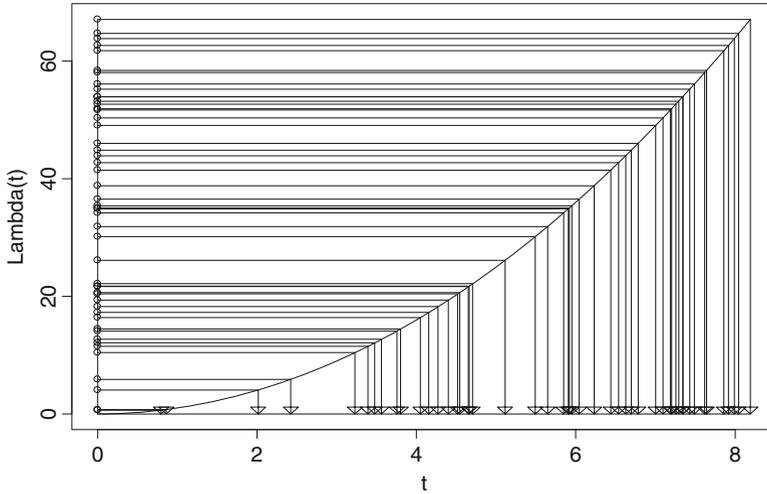


Fig. 6.4 Illustration of the inversion method when $\lambda(t) = 2t$ so that $\Lambda(t) = t^2$ and $\Lambda^{-1}(s) = \sqrt{s}$. The \circ 's on the vertical axis represent arrival times for the rate-1 base process while the ∇ 's on the horizontal axis represent arrival times in the nonstationary process

potential arrivals to form the actual arrival process. The probability that a potential arrival at time t is thinned is $1 - \lambda(t)/\tilde{\lambda}$. The algorithm is as follows:

1. Set indices $n = 1$ and $k = 1$ and $\tilde{S}_0 = 0$
2. Generate \tilde{A}_n and let $\tilde{S}_n = \tilde{S}_{n-1} + \tilde{A}_n$
3. Generate $U \sim U(0, 1)$
4. If $U \leq \lambda(\tilde{S}_n)/\tilde{\lambda}$ then
 - a. $S_k = \tilde{S}_n$
 - b. $A_k = S_k - S_{k-1}$
 - c. $k = k + 1$
- Endif
5. $n = n + 1$
6. Go to Step 2

If we again let $N(t) = \max\{n \geq 0 : S_n \leq t\}$, then Gerhardt and Nelson (2009) showed that $E(N(t)) = \Lambda(t) = \int_0^t \lambda(s) ds$. Thus, thinning \tilde{S}_n transforms the arrival times of an equilibrium renewal process with arrival rate $\tilde{\lambda} = \max_t \lambda(t)$ into a nonstationary arrival process with rate $\lambda(t) = d\Lambda(t)/dt$. While inversion and thinning both lead to arrival processes with the same time-varying arrival rate, they do not give processes that are probabilistically the same, in general (although they do for Poisson arrival processes; see Sect. 6.3.1.4). In other words, inversion and thinning are transformations that both yield the desired arrival rate, but other aspects of the arrivals, such as the variance of the number of arrivals by time t , will be different.

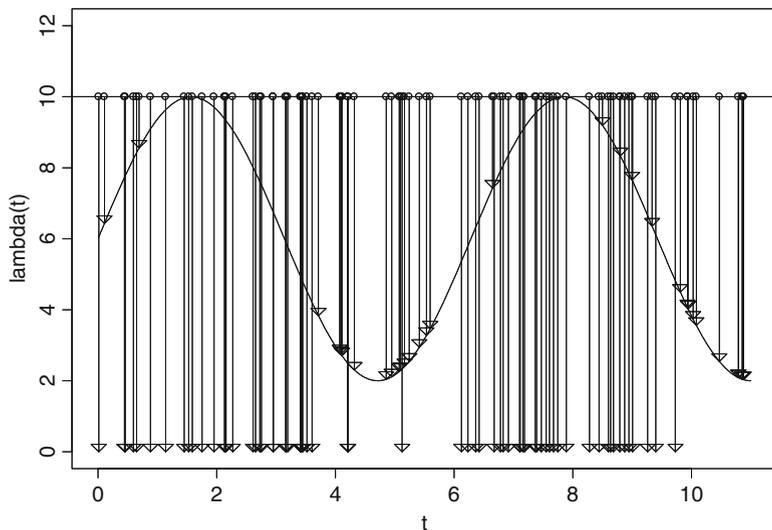


Fig. 6.5 Illustration of the thinning method when $\lambda(t) = 6 + 4\sin(t)$. The \circ 's represent arrival times for the rate-10 base process while the ∇ 's on the *horizontal axis* represent arrival times in the nonstationary process

Thinning was the approach used for the parking lot example in Sect. 3.1. Thinning has the distinct advantage of being applicable for any bounded arrival rate $\lambda(t)$, no matter how complex, while inverting $\Lambda(t)$ can be computationally difficult. However, thinning can also be slow if $\hat{\lambda}$ is greater than much of $\lambda(t)$ because many potential arrivals are generated, but most are thinned. Figure 6.5 illustrates the approach.

6.3.1.3 Estimating $\Lambda(t)$ or $\lambda(t)$ from Arrival Data

We next turn to the topic of estimating $\Lambda(t)$ or $\lambda(t)$.

Suppose that we can observe k independent realizations of an arrival process. Specifically, we observe T_{ij} , the time of the i th arrival on the j th realization for $j = 1, 2, \dots, k$; and we have reason to believe that realizations are independent and identically distributed. For instance, to model the arrival of e-mail messages to mail.iems.northwestern.edu on Mondays, we might observe $k = 10$ Mondays from 4 a.m. (time 0) to 10 p.m. (time T). For input modeling, we want to approximate the arrivals by an arrival-counting process $N(t)$ that has $E(N(t)) = \Lambda(t)$ over $0 \leq t \leq T$.

Let $C_j(t)$ be the cumulative number of arrivals by time t on the j th realization; therefore the full data set is

$$\{T_{ij}; i = 1, 2, \dots, C_j(T)\}, j = 1, 2, \dots, k. \tag{6.18}$$

The natural estimator of $\Lambda(t)$ is

$$\bar{\Lambda}(t) = \frac{1}{k} \sum_{j=1}^k C_j(t) \quad (6.19)$$

the average number of arrivals by time t over k realizations. Exercise 13 asks you to prove that $\bar{\Lambda}(t) \xrightarrow{a.s.} \Lambda(t)$ as $k \rightarrow \infty$, and $E(\bar{\Lambda}(t)) = \Lambda(t)$ for any fixed $t \in [0, T]$.

However, $\bar{\Lambda}(t)$ is not a very satisfying estimator. Let $C = \sum_{j=1}^k C_j(T)$ be the total number of observed arrivals, and let $T_{(1)} \leq T_{(2)} \leq \dots \leq T_{(C)}$ be *all* of the arrival times in (6.18) sorted from smallest to largest. Notice that $\bar{\Lambda}(t)$ is a piecewise constant function that jumps $1/k$ at each arrival time $T_{(i)}$. As a result, if we apply the inversion method to generate arrivals, then only the arrival times $T_{(i)}, i = 1, 2, \dots, C$ can be generated and nothing in between.

A solution is to linearly interpolate between observed arrival times (Leemis 1991). Let $T_{(0)} = 0$ and $T_{(C+1)} = T$. Then define

$$\hat{\Lambda}(t) = \left(\frac{C}{C+1} \right) \left\{ \frac{i}{k} + \frac{1}{k} \left(\frac{t - T_{(i)}}{T_{(i+1)} - T_{(i)}} \right) \right\} \quad (6.20)$$

when $T_{(i)} < t \leq T_{(i+1)}$ for $i = 0, 1, \dots, C$. The factor $C/(C+1)$ is needed because there are $C+1$ gaps when we include $T_{(0)}$ and $T_{(C+1)}$. Much like the linearly interpolated ecdf, Leemis (1991) showed that $\hat{\Lambda}(t) \xrightarrow{a.s.} \Lambda(t)$ as $k \rightarrow \infty$, so $\hat{\Lambda}$ fills in the gaps while still giving a consistent estimator of Λ .

To illustrate the two estimators with a (very) small example, suppose we observed an arrival process for $k = 2$ observation periods, each of $T = 5$ h. On the first observation period arrivals occurred at times $T_{11} = 1.2$ and $T_{21} = 4.1$ h, so $C_1(T) = 2$ arrivals. On the second observation period we observed only $C_2(T) = 1$ arrival at time $T_{12} = 2.4$ h. Therefore, the sorted arrival times are $T_{(1)} = 1.2, T_{(2)} = 2.4$ and $T_{(3)} = 4.1$, and $C = C_1(T) + C_2(T) = 3$.

The natural estimator $\bar{\Lambda}(t)$ jumps $1/k = 1/2$ at each arrival time $T_{(1)} = 1.2, T_{(2)} = 2.4$ and $T_{(3)} = 4.1$ and is constant in between; it is shown as the solid curve in Fig. 6.6. The interpolated estimator $\hat{\Lambda}(t)$ increases by $(C/(C+1))(1/k) = 3/8$ at each arrival time, and linearly interpolates in between; it is shown as the dashed curve in the figure. Figure 6.7 is a more realistic integrated rate function: the arrival of emergency phone calls during a 24-h period observed over three Mondays.

The algorithm below generates arrival times S_n and interarrival times A_n for a process with integrated rate function $\hat{\Lambda}$ by inversion. Remember that \tilde{S}_n and \tilde{A}_n are the arrival times and interarrival times, respectively, of an equilibrium renewal process with arrival rate 1. The algorithm is based on Leemis (1991).

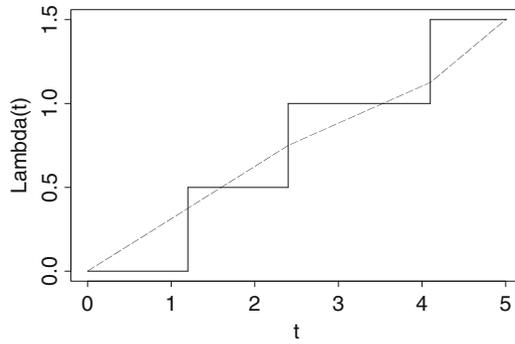


Fig. 6.6 Illustration of the estimators $\bar{\Lambda}$ (solid) and $\hat{\Lambda}$ (dashed)

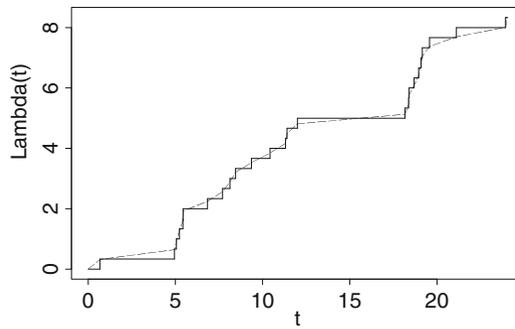


Fig. 6.7 Estimation of $\Lambda(t)$, the expected number of emergency phone calls by time t with $\bar{\Lambda}$ (solid) and $\hat{\Lambda}$ (dashed)

1. Set $n = 1$ and $A_0 = 0$
2. Generate $\tilde{S}_1 \sim G_e$
3. While $\tilde{S}_n \leq C/k$ do
 - a. $m = \left\lfloor \left(\frac{C+1}{C} \right) k\tilde{S}_n \right\rfloor$
 - b. $S_n = T_{(m)} + (T_{(m+1)} - T_{(m)}) \left(\left(\frac{C+1}{C} \right) k\tilde{S}_n - m \right)$
 - c. $A_n = S_n - S_{n-1}$
 - d. $n = n + 1$
 - e. Generate $\tilde{A}_n \sim G$
 - f. $\tilde{S}_n = \tilde{S}_{n-1} + \tilde{A}_n$
- Loop

Notice that the condition $\tilde{S}_n \leq C/k$ is needed because $\tilde{S}_n = C/k$ maps into an arrival at time T , which is the end of the observation interval.

Next consider estimating the arrival rate function $\lambda(t)$ directly. A standard method is to assume that $\lambda(t)$ is piecewise constant over intervals of length $\delta > 0$, for δ small enough. Then $\hat{\lambda}(t)$ is a piecewise constant rate function obtained as the average number of arrivals observed in nonoverlapping intervals of size δ .

To be specific, assume T/δ is integer. Then

$$\hat{\lambda}(t) = \frac{1}{k\delta} \sum_{j=1}^k [C_j(\ell(t+\delta)) - C_j(\ell(t))] \quad (6.21)$$

where $\ell(t) = \lfloor t/\delta \rfloor \delta$ is the beginning of the interval in which time t falls. This is a rather complicated way to express a simple idea: To estimate the arrival rate between, say, times $i\delta < t \leq (i+1)\delta$, compute the average number of arrivals that occurred during this interval across the k realizations, then divide by δ to make it a rate. The resulting $\hat{\lambda}(t)$ can be incorporated into a thinning algorithm, or integrated and used with inversion.

To illustrate, consider again the small example where we observed an arrival process for $k = 2$ observation periods, each of $T = 5$ h. On the first observation period arrivals occurred at times $T_{11} = 1.2$ and $T_{21} = 4.1$ h, while on the second observation period we observed only 1 arrival at time $T_{12} = 2.4$ h. If we set $\delta = 2.5$ h, then there were in total two arrivals between $t = 0 \times \delta = 0$ and $t = 1 \times \delta = 2.5$, and only one arrival between times $t = 2.5$ and $t = 2\delta = 5$. Therefore, the estimated arrival rate is

$$\lambda(t) = \begin{cases} \frac{1}{k\delta} 2 = \frac{2}{5}, & 0 < t \leq \delta = 2.5 \\ \frac{1}{k\delta} 1 = \frac{1}{5}, & 2.5 < t \leq 2\delta = 5. \end{cases}$$

Selecting δ is clearly difficult: too small and there may be intervals with few or no observed arrivals; too large and the nonstationary behavior may be masked. Henderson (2003) showed a consistency property of this estimator as $k \rightarrow \infty$ provided $\delta \rightarrow 0$.

For example, Fig. 6.8 shows $\hat{\lambda}(t)$ constructed from the data on arrivals of emergency phone calls during a 24-h period over three Mondays with $\delta = 1$ h. Notice that the sparseness of the data results in many hours of the day having estimated arrival rate 0; while the arrival rate of emergency calls may be quite small, it is unlikely to be 0 over the long run. However, if we make δ much larger then the time-of-day effect will be lost; the linearly interpolated estimator $\hat{\Lambda}(t)$ avoids this problem. To estimate the arrival rate directly it is important to have sufficient data (observation periods).

In some situations it is more convenient to collect arrival counts rather than arrival times. For instance, counts of customer arrivals might be recorded every 30 min. This makes it natural to use the piecewise constant arrival-rate function (6.21) with $\delta = 30$, the resolution used to collect the data. All else being equal, we recommend using $\hat{\Lambda}(t)$ when actual arrival times are available, and $\hat{\lambda}(t)$ when all we have are counts.

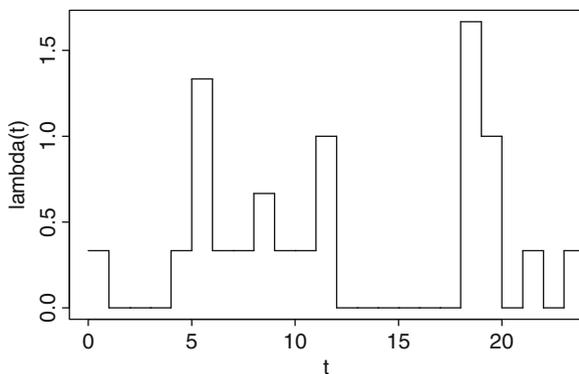


Fig. 6.8 Estimation of $\lambda(t)$, the arrival rate of emergency phone calls at time t , by $\hat{\lambda}$

6.3.1.4 Poisson or Not Poisson?

If the interarrival times of the renewal base process are exponential with rate $\tilde{\lambda}$, then the base process is Poisson, and either inversion or thinning lead to probabilistically identical *nonstationary Poisson processes (NSPP)*. The stationary and nonstationary Poisson processes are the most widely used in simulation practice—in fact, they are the default arrival process in many simulation languages—because they are often good representations of arrivals from a large population of potential arrival entities, each making independent decisions about whether or when to arrive (see Sect. 6.2.1.1). They are also very convenient processes to simulate because $G_e(t) = G(t) = 1 - e^{-\tilde{\lambda}t}, t \geq 0$, so that interarrival times of the base process are easily generated by the inverse cdf method.

How can we recognize that an arrival process is not well-represented as a NSPP, and how can we model the arrival process when it is not? Consider a NSPP with integrated rate function $\Lambda(t)$. One of the many known properties of a NSPP is that

$$\frac{\text{Var}(N(t))}{\text{E}(N(t))} = 1 \text{ for all } t \geq 0. \tag{6.22}$$

In words, the ratio of the variance of the number of arrivals to the expected number of arrivals by time t is always 1, no matter what $\Lambda(t)$ or $\lambda(t)$ is chosen. However, some real arrival processes are more variable (e.g., call centers) or less variable (e.g., manufacturing orders) than this, which is one way to recognize and adjust for departures from Poisson.

Gerhardt and Nelson (2009) showed that if the nonstationary arrival process is generated by the inversion method then

$$\frac{\text{Var}(N(t))}{\text{E}(N(t))} \approx \sigma_A^2 \text{ for large } t \tag{6.23}$$

Table 6.1 Illustration of estimating $\hat{\sigma}_A^2$ given hourly arrival count data from $k = 3$ realizations. The # column is the hourly count and $C_j(t)$ is the cumulative count

t	#	$C_1(t)$	#	$C_2(t)$	#	$C_3(t)$	$\bar{\Lambda}(t)$	$V(t)$	$V(t)/\bar{\Lambda}(t)$
8 a.m.–9 a.m.	1	1	3	3	1	1	1.67	1.33	0.80
9 a.m.–10 a.m.	5	6	7	10	4	5	7.00	7.00	1.00
10 a.m.–11 a.m.	14	20	8	18	11	16	18.00	4.00	0.22
11 a.m.–12 p.m.	19	39	15	33	17	33	35.00	12.00	0.34
12 p.m.–1 p.m.	22	61	20	53	18	51	55.00	28.00	0.51
1 p.m.–2 p.m.	9	70	11	64	8	59	64.33	30.33	0.47
2 p.m.–3 p.m.	4	74	1	65	2	61	66.67	44.33	0.67
3 p.m.–4 p.m.	2	76	3	68	1	62	68.67	49.33	0.72
$\hat{\sigma}_A^2$									0.59

where $\sigma_A^2 = \text{Var}(\bar{A}_2)$, the variance of the stationary interarrival times of the rate-1 base process. When the base process is exponential, then $\sigma_A^2 = 1$. Thus, the base process provides a way to obtain arrival processes that are more or less variable than a NSPP.

The key idea is that deviation of the ratio $\text{Var}(N(t))/E(N(t))$ from 1 is a measure of deviation from being a Poisson process, but the deviation can be matched by using an appropriate base process.

Assuming we have arrival data, let

$$V(t) = \frac{1}{k-1} \sum_{j=1}^k (C_j(t) - \bar{\Lambda}(t))^2$$

be the estimated variance of the number of arrivals by time t . If we select a collection of times $t_1 < t_2 < \dots < t_m$, then an estimator of σ_A^2 is

$$\hat{\sigma}_A^2 = \frac{1}{m} \sum_{i=1}^m \frac{V(t_i)}{\bar{\Lambda}(t_i)}.$$

When this value is significantly different from 1, then it indicates that a NSPP is not appropriate; it also provides an estimate of what the variance of the equilibrium renewal base process should be. In the exercises we suggest some base processes that have controllable variances and for which generating values from G and G_e is not difficult.

Estimating $\hat{\sigma}_A^2$ from data is illustrated in Table 6.1. Notice that the hourly count data must first be transformed into cumulative counts before computing $\bar{\Lambda}(t)$, $V(t)$ and their ratio $V(t)/\bar{\Lambda}(t)$. The average of these ratios is $\hat{\sigma}_A^2 = 0.59$, indicating an arrival process that is less variable than Poisson. Therefore, to generate arrivals by inversion we want a renewal base process that is correspondingly less variable than Poisson.

At present there is no easy-to-use relationship between the variance of the base process and the arrival process for thinning, so we only recommend thinning for generating a NSPP.

6.3.2 Random Vectors

Suppose that the simulation requires generation of a random vector (X_1, X_2, \dots, X_k) , where the i th component has cdf $F_i = F_{X_i}$. Let (U_1, U_2, \dots, U_k) be a vector of $U(0, 1)$ random variables that may or may not be dependent. Then it follows from the development of the inverse cdf method in Sect. 2.2.1 that the transformation

$$\begin{pmatrix} X_1 \\ X_2 \\ \vdots \\ X_k \end{pmatrix} = \begin{pmatrix} F_1^{-1}(U_1) \\ F_2^{-1}(U_2) \\ \vdots \\ F_k^{-1}(U_k) \end{pmatrix}$$

produces a random vector with the correct marginal distributions, and it seems plausible that the vector (X_1, X_2, \dots, X_k) will inherit some of the dependence of (U_1, U_2, \dots, U_k) . Therefore, the key to this method is finding a dependent vector (U_1, U_2, \dots, U_k) that produces the desired dependence among (X_1, X_2, \dots, X_k) . Here we will measure dependence by the correlation matrix $\mathbf{R} = (\rho_{ij})$ where $\rho_{ij} = \text{Corr}(X_i, X_j)$.

While there are a number of ways to construct a base vector (U_1, U_2, \dots, U_k) , we will describe the normal-to-anything (NORTA) method:

$$\begin{pmatrix} X_1 \\ X_2 \\ \vdots \\ X_k \end{pmatrix} = \begin{pmatrix} F_1^{-1}(U_1) \\ F_2^{-1}(U_2) \\ \vdots \\ F_k^{-1}(U_k) \end{pmatrix} = \begin{pmatrix} F_1^{-1}[\Phi(Z_1)] \\ F_2^{-1}[\Phi(Z_2)] \\ \vdots \\ F_k^{-1}[\Phi(Z_k)] \end{pmatrix} \quad (6.24)$$

where (Z_1, Z_2, \dots, Z_k) has a standard multivariate normal distribution with correlation matrix \mathbf{R} , each component having mean 0 and variance 1, and Φ is the cdf of the standard normal distribution. The NORTA method is based on the fact that if Z is standard normal, then $U = \Phi(Z)$ has a $U(0, 1)$ distribution; see Exercise 36. There are many methods for constructing joint distributions of uniforms (U_1, U_2, \dots, U_k) ; these distributions are known as *copulas*. The normal copula is a good one for controlling the correlations of (X_1, X_2, \dots, X_k) , but other measures of dependence may be relevant in some applications. A comprehensive reference is Biller and Corlu (2012).

To employ the NORTA method, we first choose or fit the cdfs F_1, F_2, \dots, F_k of the k component random variables, and also estimate the pairwise correlations $\mathbf{R} = (\rho_{ij})$ that we require. Given F_1, F_2, \dots, F_k and ρ_{ij} , we find a correlation matrix $\mathbf{R} = (r_{ij})$ for (Z_1, Z_2, \dots, Z_k) that implies the target correlation matrix \mathbf{R} for (X_1, X_2, \dots, X_k) after the NORTA transformation. Variate generation is done by first generating a standard multivariate-normal random vector (Z_1, Z_2, \dots, Z_k) with correlation matrix \mathbf{R} and then applying the transformation (6.24).

The level of complexity involved in implementing the NORTA method is significantly greater than other methods described here. Details are provided in the

Appendix to the chapter. To illustrate we use a very simple example chosen to be easy to understand, rather than corresponding to a practical problem.

Suppose that we need bivariate random vectors (X_1, X_2) where X_1 has an exponential distribution with mean 10, X_2 has a discrete uniform distribution on $\{1, 2, \dots, 10\}$, and we want to match some $\rho_{12} = \text{Corr}(X_1, X_2)$. These might represent a time to do something and a count of the number of things that are done, respectively, so they are naturally dependent. The NORTA transformation is

$$\begin{pmatrix} X_1 \\ X_2 \end{pmatrix} = \begin{pmatrix} F_1^{-1}(U_1) \\ F_2^{-1}(U_2) \end{pmatrix} = \begin{pmatrix} -10\ln(1-U_1) \\ \lceil 10U_2 \rceil \end{pmatrix} = \begin{pmatrix} -10\ln[1-\Phi(Z_1)] \\ \lceil 10\Phi(Z_2) \rceil \end{pmatrix}$$

where (Z_1, Z_2) have a standard bivariate normal distribution with correlation r_{12} . The key to the NORTA method is finding r_{12} that implies (X_1, X_2) have correlation ρ_{12} after the transformation.

Using the method described in the Appendix to the chapter we solved the correlation-matching problem for several possible values of ρ_{12} as shown below. Notice that the relationship is not linear. Because of the symmetry of the discrete uniform distribution in this problem we can obtain correlation $-\rho_{12}$ by using base correlation $-r_{12}$; this is not the case for all pairs of marginal distributions.

ρ_{12}	r_{12}
-0.5	-0.578
0.5	0.578
0.7	0.813
0.8	0.950
0.85	0.991

Figure 6.9 shows a scatter plot of 200 (X_1, X_2) pairs when the correlation is $\rho_{12} = 0.7$. Notice how large values of X_1 tend to occur with large values of X_2 .

6.4 Generating Random Variates

The inverse cdf method of random-variate generation was described in Sect. 2.2.1. For generating univariate inputs this is, in theory, the only method we need. The fact that inversion maps a single uniform U monotonically into a single X via $X = F_X^{-1}(U)$ will be helpful for experiment design, as discussed in Chap. 8, so inversion is the method of choice.

In practice, however, each inversion may require solving the root-finding problem

$$U = F_X(X) = \int_{-\infty}^X f_X(x) dx \quad (6.25)$$

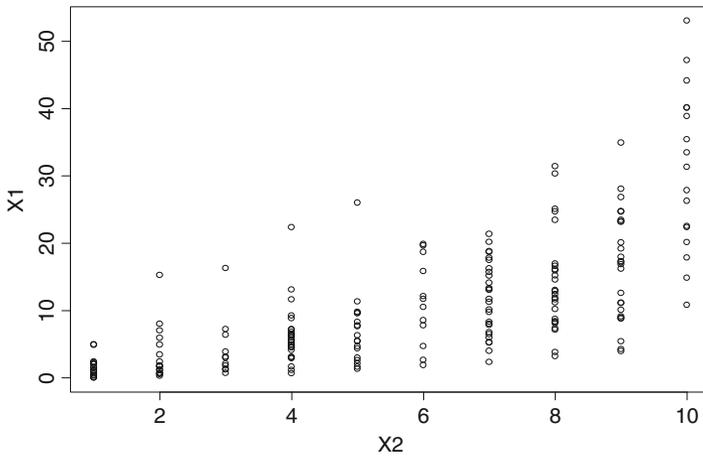


Fig. 6.9 Scatter plot of 200 pairs of bivariate exponential (X_1) and discrete uniform (X_2) variates with correlation 0.7

for X , which could be slow if numerical integration is required. In the case of a discrete distribution, a search over a countably infinite number of possible outcomes might be necessary (e.g., the Poisson distribution). Therefore, methods other than inversion are sometimes more practical. Nevertheless, it is worth noting that even when there is no closed-form expression for F_X^{-1} , if (6.25) can be solved efficiently and to sufficient numerical accuracy, then inversion is still preferred.

Random-variate generation is a vast topic with decades of development. An outstanding general reference is Devroye (1986), along with the update in Devroye (2006). The most general method, after inversion, is rejection, which we cover in Sect. 6.4.1. Most other approaches exploit particular properties of the distribution from which variates are needed; we provide a few example in Sect. 6.4.2.

6.4.1 Rejection Method

Let X be the random variable of interest. The rejection method applies when we can express

$$\Pr\{X \leq x\} = \Pr\{V \leq x | \mathcal{A}\}$$

where V is another random variable and \mathcal{A} is some “acceptance” event (for this reason the method is sometimes called “acceptance–rejection”). If V is easy to generate, while X is difficult, then this suggests that we can generate X in the following way:

1. Generate V
2. If \mathcal{A} occurs, then return $X = V$
 Otherwise, reject V and go to Step 1

Provided generation of V is fast and, on average, we do not have to reject too many V 's until the event \mathcal{A} occurs, then rejection can be a competitive method.

How can we set up such a situation? Suppose that X has a density function f_X (and cdf F_X), and let $m(x)$ be a function that “majorizes” f_X , by which we mean $m(x) \geq f_X(x)$ for all x . Unless they are equal, $m(x)$ will not be a density function (its integral will be greater than 1), but

$$g(x) = \frac{m(x)}{\int_{-\infty}^{\infty} m(y) dy} = \frac{m(x)}{c}$$

will be. The following is a general rejection algorithm:

1. Generate $V \sim g$
2. Generate $U \sim U(0, 1)$
3. If $U \leq f_X(V)/m(V)$ then return $X = V$
Otherwise go to Step 1

Here the acceptance event is $\mathcal{A} = \{U \leq f_X(V)/m(V)\}$, and we need to show that

$$\Pr\{V \leq x | \mathcal{A}\} = F_X(x).$$

By definition

$$\Pr\{V \leq x | \mathcal{A}\} = \frac{\Pr\{V \leq x, \mathcal{A}\}}{\Pr\{\mathcal{A}\}}.$$

But

$$\begin{aligned} \Pr\{\mathcal{A}\} &= \Pr\{U \leq f_X(V)/m(V)\} \\ &= \int_{-\infty}^{\infty} \Pr\{U \leq f_X(y)/m(y) | V = y\} g(y) dy \\ &= \int_{-\infty}^{\infty} \frac{f_X(y)}{m(y)} \cdot \frac{m(y)}{c} dy \\ &= \frac{1}{c} \int_{-\infty}^{\infty} f_X(y) dy = \frac{1}{c}. \end{aligned} \tag{6.26}$$

A similar argument gives the numerator:

$$\begin{aligned} \Pr\{V \leq x, U \leq f_X(V)/m(V)\} &= \int_{-\infty}^{\infty} \Pr\{y \leq x, U \leq f_X(y)/m(y) | V = y\} g(y) dy \\ &= \int_{-\infty}^x \Pr\{U \leq f_X(y)/m(y)\} g(y) dy \\ &= \frac{1}{c} \int_{-\infty}^x f_X(y) dy = \frac{1}{c} F_X(x). \end{aligned} \tag{6.27}$$

Combining (6.27) and (6.26) proves the result.

Notice that (6.26) is the probability that we accept V on any trial; since the trials are independent, they have a geometric distribution and the expected number of trials to generate one X is c . Thus, the closer c is to 1 the more efficient the algorithm is.

Also notice that the rejection method requires a random number of $U(0, 1)$'s to generate each X .

As an example, consider the beta distribution with pdf

$$f_X(x) = \frac{x^{\alpha_1-1}(1-x)^{\alpha_2-1}}{B(\alpha_1, \alpha_2)} \tag{6.28}$$

for $0 \leq x \leq 1$, where $\alpha_1, \alpha_2 > 0$ are shape parameters, and

$$B(\alpha_1, \alpha_2) = \int_0^1 u^{\alpha_1-1}(1-u)^{\alpha_2-1} du.$$

The beta distribution does not have a closed-form inverse cdf.

When $\alpha_1 > 1$ and $\alpha_2 > 1$ the mode of the beta distribution occurs at $x^* = (\alpha_1 - 1)/(\alpha_1 + \alpha_2 - 2)$. Thus, the highest point on the density function is $f^* = f_X(x^*)$ and the constant function

$$m(x) = f^*, \quad 0 \leq x \leq 1$$

majorizes f_X . This implies that $g(x) = 1$, for $0 \leq x \leq 1$, the uniform distribution on $(0, 1)$, This gives the following rejection algorithm for generating beta random variates when $\alpha_1 > 1$ and $\alpha_2 > 1$:

0. Compute $f^* = f_X\left(\frac{\alpha_1-1}{\alpha_1+\alpha_2-2}\right)$
1. Generate $V \sim g = U(0, 1)$
2. Generate $U \sim U(0, 1)$
3. If $U \leq f_X(V)/m(V) = [V^{\alpha_1-1}(1-V)^{\alpha_2-1}/B(\alpha_1, \alpha_2)]/f^*$ then return $X = V$
- Otherwise go to Step 1

The efficiency of this rejection algorithm depends on the particular values of α_1 and α_2 . For instance, when $\alpha_1 = 4$ and $\alpha_2 = 3$, then $c = 2.0736$ meaning the expected number of trials is a little greater than 2 to generate each beta variate. By modern standards this is not a very efficient rejection algorithm. This particular example is illustrated in Fig. 6.10.

State-of-the-art algorithms use a number of tricks to increase efficiency. One is to express the density f_X as a probabilistic mixture of pieces so that each piece can be majorized very closely.

6.4.2 Particular Properties

By ‘‘particular properties’’ we mean relationships among distributions that facilitate variate generation. For instance, if $X_1 \sim \text{gamma}(\alpha_1, \beta)$, independent of $X_2 \sim \text{gamma}(\alpha_2, \beta)$, then $X = X_1/(X_1 + X_2)$ has a beta(α_1, α_2) distribution. Thus, if one has a gamma random-variate-generation algorithm, then it can be used to generate beta variates.

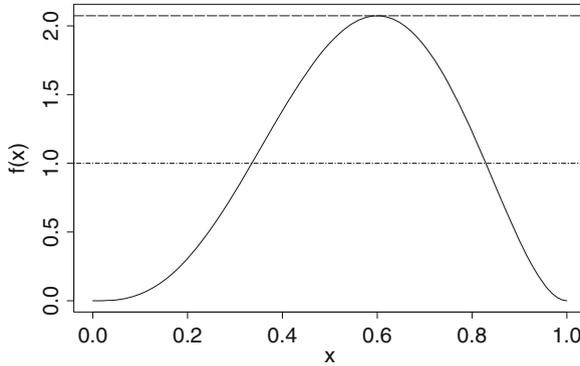


Fig. 6.10 Rejection method for a beta distribution (solid line), with majorizing function $m(x)$ (dashed line) and corresponding density $g(x)$ (dot-dashed line)

Relationships among the standard normal distribution, the normal distribution, and the lognormal distribution have already been mentioned: If $Z \sim N(0, 1)$, then $X = \mu + \sigma Z$ has a $N(\mu, \sigma^2)$ distribution. Further, $Y = \exp(X)$ has a lognormal distribution. A good reference for such relationships is Leemis and McQueston (2008).

Although particular properties can often supply competitive random-variate generation algorithms, they should typically be used only if they provide exact methods. Recall that what we mean by “exact” is that if we had infinite-precision computer arithmetic, and a source of truly $U(0, 1)$ random numbers, then $\Pr\{X \leq x\} = F_X(x)$. There are particular properties that only lead to approximate methods. For instance, if X has a Poisson distribution with mean λ and λ is large, then $(X - \lambda)/\sqrt{\lambda}$ is approximately standard normal. Thus, if we had a method for generating $N(\lambda, \lambda)$ random variates, then we could round them to obtain approximately Poisson distributed variates. However, there is no reason to settle for this approximation since exact, uniformly fast methods exist (e.g., Hörmann 1993).

The most infamous example of an approximate method is generating a $N(0, 1)$ random variate by setting

$$Z = \sum_{i=1}^{12} U_i - 6 \quad (6.29)$$

where U_1, U_2, \dots, U_{12} are i.i.d. $U(0, 1)$. The justification for normality is the central limit theorem, and the convenient choice of 12 gives variance 1. In Exercise 24 you are asked to explore the faults of this method.

6.5 Generating Pseudorandom Numbers

The driving source of randomness in stochastic simulations is, in theory, i.i.d. samples from the $U(0, 1)$ distribution; but in practice they are a deterministic sequence

of pseudorandom numbers that appear to be random. As described in Sect. 2.2.2, one way to visualize the pseudorandom numbers is as a large, ordered list

$$u_1, u_2, u_3, \dots, u_i, u_{i+1}, \dots, u_{P-1}, u_P, u_1, u_2, u_3, \dots$$

where P is the period. Here we use lowercase u to denote pseudorandom numbers to emphasize that they are deterministic and not actually random.

Rather than storing an actual list, pseudorandom numbers are invariably generated by a recursive algorithm called a *pseudorandom-number generator*. In this section we present enough background about pseudorandom-number generators to use them effectively. It is worth stating that inventing algorithms to generate a sequence that appears to be i.i.d. $U(0, 1)$ is quite difficult. Basic properties, such as having the right mean ($1/2$) and variance ($1/12$), are relatively easy to achieve, but apparent independence is not. For instance, given any dimension d , the values $(u_{i+1}, u_{i+2}, \dots, u_{i+d}), i = 1, 2, \dots$ should appear to be uniformly distributed in the d -dimensional unit hypercube $(0, 1)^d$. This property is hard to achieve as d increases, and poor generators can give results that do not look random in high dimensions. Creating pseudorandom-number generators that have a long period P and good statistical properties involves number theory, computer science and clever testing. A very accessible reference is L'Ecuyer (2006), to which the reader should refer for technical details not covered here.

The pseudorandom-number generators we consider produce a sequence of integer values that are scaled to be between 0 and 1. The generator in VBASim starts with an integer seed $z_0 \in \{1, 2, \dots, 2^{31} - 1\}$, and then generates pseudorandom numbers via the recursion

$$\begin{aligned} z_i &= 630,360,016z_{i-1} \bmod 2^{31} - 1 \\ u_i &= \frac{z_i}{2^{31} - 1}. \end{aligned} \tag{6.30}$$

This is a VBA implementation of the algorithm described in Marse and Roberts (1983) and Law (2007). Recall that “mod m ” means to retain the remainder after dividing by m ; for instance, $7 \bmod 3 = 1$. This generator is one instance of a class of generators of the form

$$\begin{aligned} z_i &= az_{i-1} \bmod m \\ u_i &= \frac{z_i}{m} \end{aligned} \tag{6.31}$$

which are called (linear) multiplicative congruential generators (MCGs). MCGs have been in widespread use for many years. Here are some important observations about MCGs:

1. The only possible values a MCG can return are $\{0, 1/m, 2/m, \dots, (m-1)/m\}$. Clearly we do not want it to return 0 since it will produce nothing but zeroes

from then on.⁵ Thus, the maximal period of a MCG is $P = m - 1$, which the generator (6.30) achieves.

2. We want m to be large so that the interval between 0 and 1 is densely filled; thus, $m = 2^{31} - 1$ is a natural choice because it is the largest integer representable on 32-bit computers. However, to appear random requires more than just a long period; there are 534 million values of a that achieve full period for $m = 2^{31} - 1$, but only a fraction of them generate sequences with good statistical properties.
3. The computation $z_i = az_{i-1} \bmod m$ must be done with care. For example, the multiplication $630,360,016z_{i-1}$ will frequently yield a value greater than $2^{31} - 1$, leading to overflow if we try to do integer arithmetic.
4. Exercise 45 asks you to show that $z_i = a^i z_0 \bmod m$ for a MCG. Therefore, it is possible to skip ahead in the sequence since

$$z_i = a^i z_0 \bmod m = (a^i \bmod m) z_0 \bmod m.$$

Thus, a single generator with a long period can act as several virtual generators, each one starting with their own unique “seed” z_0 spaced far apart. The VBASim implementation of the generator (6.30) has 100 virtual generators with seeds spaced 100,000 pseudorandom numbers apart; these are accessed via the random-number streams 1–100. For instance, stream 1 corresponds to $z_0 = 1,973,272,912$ while stream 2 has $z_0 = 281,629,770$.

The pseudorandom-number generator (6.30) is adequate for learning about stochastic simulation, but with a period of only about two billion it is insufficient for many problems. Consider this simple and very realistic example: We want to estimate the reliability of a system for which failure needs to be very unlikely because human life is involved. Suppose each replication of the simulation requires 100 random numbers, and we would like to observe at least 30 failures. If the actual probability of failure is one in a million, then we would expect to need $100 \times 30 \times 1,000,000 = 3,000,000,000$ random numbers, many more than the pseudorandom-number generator (6.30) can provide before repeating.

Unfortunately, even if the period P of a generator is sufficient for the applications we have in mind problems can occur. L’Ecuyer and Simard (2001) demonstrated that for many standard generators, such as (6.30), nonrandom behavior of the generator is easily detectable by statistical tests if more than a fraction of the period is actually used (from $P^{1/3}$ to $P^{1/2}$). This implies that we would like to have generators with periods substantially larger than what the application requires. To complete our understanding of modern pseudorandom number generators we describe two general approaches that greatly extend the period of MCGs; selecting the specific constants that give good statistical properties is, however, beyond the scope of this book.

⁵ Typically we do not want pseudorandom-number generators to return either 0 or 1 because these values correspond to $\pm\infty$ for some distributions from which we want to generate random variates. For instance, the inverse cdf of the exponential distribution with mean 1 is $F^{-1}(u) = -\ln(1 - u)$.

6.5.1 Multiple Recursive Generators

A recursive pseudorandom-number generator maps a previous state into a new state, so its period depends on the number of distinct states it can take. The period of a MCG is limited because its state is the most recently generated integer z_{i-1} , which can assume only m values. A multiple recursive generator (MRG) of order K extends the state to include the K most recent values:

$$z_i = (a_1 z_{i-1} + a_2 z_{i-2} + \cdots + a_K z_{i-K}) \bmod m$$

$$u_i = \begin{cases} \frac{z_i}{m+1}, & z_i > 0 \\ \frac{m}{m+1}, & \text{otherwise.} \end{cases} \quad (6.32)$$

Notice that a value of $z_i = 0$ is not fatal for a MRG, but is mapped into $m/(m+1)$ to avoid generating $u_i = 0$. Since the state of this generator is $(z_{i-K}, z_{i-K+1}, \dots, z_{i-1})$, and since each element can take m possible values, the maximal period—avoiding the state $(0, 0, \dots, 0)$ —is $P = m^K - 1$. A necessary (but not sufficient) condition for achieving full period is that a_K and at least one other a_j are not zero. Notice that the “seed” for such a generator consists of K values $(z_0, z_1, \dots, z_{K-1})$.

6.5.2 Combined Generators

Combining MRGs is a way to further increase the state of the generator. Suppose that we have $j = 1, 2, \dots, J$ MRGs of the form

$$z_{i,j} = (a_{1,j} z_{i-1,j} + a_{2,j} z_{i-2,j} + \cdots + a_{K,j} z_{i-K,j}) \bmod m_j.$$

That is, $z_{i,j}$ is the i th pseudorandom integer from the j th MRG. Then we form the i th pseudorandom number from

$$z_i = (\delta_1 z_{i,1} + \delta_2 z_{i,2} + \cdots + \delta_J z_{i,J}) \bmod m_1$$

$$u_i = \begin{cases} \frac{z_i}{m_1+1}, & z_i > 0 \\ \frac{m_1}{m_1+1}, & \text{otherwise} \end{cases}$$

where $\delta_1, \delta_2, \dots, \delta_J$ are fixed integers. In words, the generator consists of J MRGs of order K executed in parallel, with their outputs combined to form a pseudorandom number.

If each of the component MRGs uses a prime number modulus m_j and has period $P_j = m_j^K - 1$, then the period of the combined generator is at most $P =$

$P_1 P_2 \cdots P_k / 2^{J-1}$. Notice that the initial “seed” for such a generator consists of $J \times K$ values.

A specific instance that has been widely implemented is MRG32k3a (L’Ecuyer 1999), which consists of $J = 2$ MRGs each of order $K = 3$:

$$\begin{aligned} z_{i,1} &= (1,403,580 z_{i-2,1} - 810,728 z_{i-3,1}) \bmod 2^{32} - 209 \\ z_{i,2} &= (527,612 z_{i-1,2} - 1,370,589 z_{i-3,2}) \bmod 2^{32} - 22,853 \\ z_i &= (z_{i,1} - z_{i,2}) \bmod 2^{32} - 209. \end{aligned}$$

The period of this generator is $P \approx 2^{191} \approx 3 \times 10^{57}$, which is enormous: If you could generate two billion pseudorandom numbers per second (the entire period of the VBA_{Sim} generator), then it would take longer than the age of the universe to exhaust the period of MRG32k3a!

There are reasons why combining MRGs can lead to a generator with good statistical properties as well as long period, but again the specific choices must be made with great care; see L’Ecuyer (2006). The following result shows that if the first component MRG $z_{i,1}$ provides what appear to be uniformly distributed integers, then the combined generator will as well. Thus, combining several MRGs will not destroy the uniformity of the first generator. Of course, this provides no guarantee that successive values appear to be independent.

Theorem 6.2 (L’Ecuyer 1988). *Suppose that Z_1, Z_2, \dots, Z_J are independent, integer-valued random variables, and that Z_1 has a discrete uniform distribution on $\{0, 1, \dots, m_1 - 1\}$. Then*

$$Z = \sum_{j=1}^J Z_j \bmod m_1$$

also has a discrete uniform distribution on $\{0, 1, \dots, m_1 - 1\}$.

Proof. Let $W = \sum_{j=2}^J Z_j \bmod m_1$, and notice that $Z_1 \bmod m_1 = Z_1$. Then

$$\begin{aligned} Z &= \left(Z_1 \bmod m_1 + \sum_{j=2}^J Z_j \bmod m_1 \right) \bmod m_1 \\ &= (Z_1 + W) \bmod m_1. \end{aligned}$$

Then since W only takes values in $\{0, 1, \dots, m_1 - 1\}$ we have for $z = 0, 1, \dots, m_1 - 1$ that

$$\begin{aligned} \Pr\{Z = z\} &= \sum_{w=0}^{m_1-1} \Pr\{Z = z|W = w\} \Pr\{W = w\} \\ &= \sum_{w=0}^{m_1-1} \Pr\{(Z_1 + W) \bmod m_1 = z|W = w\} \Pr\{W = w\} \end{aligned}$$

$$\begin{aligned}
&= \sum_{w=0}^{m_1-1} \Pr\{(Z_1 + w) \bmod m_1 = z\} \Pr\{W = w\} \\
&= \sum_{w=0}^{m_1-1} \frac{1}{m_1} \Pr\{W = w\} \\
&= \frac{1}{m_1}
\end{aligned} \tag{6.33}$$

where (6.33) follows because there is only one value of Z_1 that will make $(Z_1 + w) \bmod m_1 = z$, and Z_1 has a discrete uniform distribution so all values are equally likely. \square

6.5.3 Proper Use of Pseudorandom-Number Generators

Pseudorandom-number generators, such as `lcgrand` in `VBASim`, are implemented so that after each call they retain memory of the generator's state. Thus, each call to the generator advances to the next pseudorandom number in the sequence, approximating an independent sample from the $U(0, 1)$ distribution. This is why we may treat replications as (pseudo) independent, as long as we do not exhaust the period P of the generator or mistakenly reset the initial seed between replications. Notice that a MCG will generate a distinct integer z_i on each call, while a MRG or combined MRGs can (and will) generate the same z_i value many times before exhausting its period, since different generator states map into the same z_i value.

For a well-tested pseudorandom-number generator such as `MRG32k3a` there is no advantage to using any particular initial seed. A common misconception is that pseudorandom-number generators include different streams (initial seeds) because some parts of the sequence are better (more random) than others, or that assigning different streams makes different input processes more independent; this is not the case. A good pseudorandom-number generator appears to produce a sequence of i.i.d. $U(0, 1)$'s, so it does not matter if they are used in their natural order or if they are obtained from different subsequences.

Of course, surprising events occur in a pseudorandom sequence just as they would in a truly random process. This is why adequate run lengths or numbers of replications are essential to overcome the quirks of randomness, and why statistical error should be measured by a standard error or a confidence interval. These are the goals of good experiment design and analysis, the subjects of Chaps. 7 and 8.

Streams exist to synchronize how the random numbers are used when we compare alternative scenarios via simulation. Roughly speaking, we want each scenario to see the same source of randomness so that observed differences in performance are due to structural differences in the scenarios, not different random numbers. As discussed in Chap. 8, this is facilitated by insuring (as far as possible) that each pseudorandom number is used for the same purpose in the simulation of each alternative

scenario. Assigning distinct streams to distinct input processes facilitates this. If only a single scenario is simulated then streams are irrelevant.

Finally, it is sometimes argued that to appropriately represent the uncertainty that occurs in reality, it is important to randomly select the initial seed z_0 . While this may appear correct on the surface, remember that the goal of a stochastic simulation is not to produce a random result, as in a game of cards or in a lottery, but rather to estimate underlying properties of a system that is subject to uncertainty. Whether the initial seed is fixed or random, the experiment should be conducted so as to overcome the effect of randomness and obtain precise estimates. If the pseudorandom-number generator can be trusted, then the benefits of repeatability and precise comparisons that can be obtained with fixed seeds or streams outweigh the superficial connection to true uncertainty provided by choosing a random seed.

Appendix: Properties of the GLD

The results below can be found in Karian and Dudewicz (2000) or Lakhany and Mausser (2000).

The pdf of the GLD is

$$f(x) = \frac{\lambda_2}{\lambda_3 y^{\lambda_3 - 1} + \lambda_4 (1 - y)^{\lambda_4 - 1}}$$

at the point $x = Q(y)$, where Q is the inverse cdf. This expression is useful for plotting a fitted GLD distribution. The support of the GLD depends on its parameters (Karian and Dudewicz 2000):

λ_3	λ_4	Support
$\lambda_3 > 0$	$\lambda_4 > 0$	$[\lambda_1 - 1/\lambda_2, \lambda_1 + 1/\lambda_2]$
$\lambda_3 > 0$	$\lambda_4 = 0$	$[\lambda_1, \lambda_1 + 1/\lambda_2]$
$\lambda_3 = 0$	$\lambda_4 > 0$	$[\lambda_1 - 1/\lambda_2, \lambda_1]$
$\lambda_3 < 0$	$\lambda_4 < 0$	$(-\infty, \infty)$
$\lambda_3 < 0$	$\lambda_4 = 0$	$(-\infty, \lambda_1 + 1/\lambda_2]$
$\lambda_3 = 0$	$\lambda_4 < 0$	$[\lambda_1 - 1/\lambda_2, \infty)$

To fit a GLD via the method of moments, the following are valid provided $\lambda_3 > -1/4$ and $\lambda_4 > -1/4$:

$$\mu = \lambda_1 + \frac{A}{\lambda_2}$$

$$\sigma^2 = \frac{B - A^2}{\lambda_2^2}$$

$$\alpha_3 = \frac{C - 3AB + 2A^3}{\lambda_2^3 \sigma^3}$$

$$\alpha_4 = \frac{D - 4AC + 6A^2B - 3A^4}{\lambda_2^4 \sigma^4}$$

where

$$A = \frac{1}{1 + \lambda_3} - \frac{1}{1 + \lambda_4}$$

$$B = \frac{1}{1 + 2\lambda_3} + \frac{1}{1 + 2\lambda_4} - 2B(1 + \lambda_3, 1 + \lambda_4)$$

$$C = \frac{1}{1 + 3\lambda_3} - \frac{1}{1 + 3\lambda_4} - 3B(1 + 2\lambda_3, 1 + \lambda_4) + 3B(1 + \lambda_3, 1 + 2\lambda_4)$$

$$D = \frac{1}{1 + 4\lambda_3} + \frac{1}{1 + 4\lambda_4} - 4B(1 + 3\lambda_3, 1 + \lambda_4) + 6B(1 + 2\lambda_3, 1 + 2\lambda_4) - 4B(1 + \lambda_3, 1 + 3\lambda_4)$$

with

$$B(a, b) = \int_0^1 x^{a-1} (1-x)^{b-1} dx.$$

Software can be found in Karian and Dudewicz (2000).

To fit via least squares we need $E[Q(U_{(i)})]$, which can be computed using the following:

$$E[U_{(i)}^{\lambda_3}] = \frac{\Gamma(m+1)\Gamma(i+\lambda_3)}{\Gamma(i)\Gamma(m+\lambda_3+1)}$$

$$E[(1-U_{(i)})^{\lambda_4}] = \frac{\Gamma(m+1)\Gamma(m-i+\lambda_4+1)}{\Gamma(m-i+1)\Gamma(m+\lambda_4+1)}.$$

Appendix: Implementation of the NORTA Method

In this appendix we provide details on one implementation of the NORTA method. We assume that you start with observed input data, fit the NORTA input model, and then generate variates.

Fitting NORTA Components

Suppose we have observed input data that we believe are i.i.d. vectors $\mathbf{X}_h = (X_{1h}, X_{2h}, \dots, X_{kh})^\top$, $h = 1, 2, \dots, m$. Since the random vectors are i.i.d., then so are $X_{i1}, X_{i2}, \dots, X_{im}$ for any one of the marginal distributions i . Thus, the i th marginal

may be fit by using any of the methods described in Sect. 6.2. Denote the fitted marginals as $\hat{F}_i, i = 1, 2, \dots, k$.

Let

$$\bar{\mathbf{X}} = \frac{1}{m} \sum_{h=1}^m \mathbf{X}_h$$

be the sample mean vector, and

$$\hat{\mathbf{C}} = \frac{1}{m-1} \sum_{h=1}^m (\mathbf{X}_h - \bar{\mathbf{X}})(\mathbf{X}_h - \bar{\mathbf{X}})^\top \quad (6.34)$$

be the sample variance-covariance matrix. The sample correlations are therefore $\hat{\rho}_{ij} = \hat{C}_{ij} / \sqrt{\hat{C}_{ii}\hat{C}_{jj}}$. Let $\hat{\mathbf{R}} = (\hat{\rho}_{ij})$ be the estimated correlation matrix. The fitted marginals $\hat{F}_i, i = 1, 2, \dots, k$ and correlations $\hat{\rho}_{ij}$ are the inputs to the NORTA correlation-matching problem; we describe one solution method below.

Under the set-up described above, $\hat{\mathbf{R}}$ is guaranteed to be positive semi-definite, a requirement to be a legitimate correlation matrix. Unfortunately, if we do not have m complete vectors of observations on all k input processes, then there is no such guarantee. For instance, (X_1, X_2) might have been observed together 100 times, while a separate study observed (X_2, X_3) together 87 times, but (X_1, X_3) might never have been observed jointly so a guess of their correlation will be used. From these data we can estimate ρ_{12}, ρ_{13} and ρ_{23} separately, but there is no longer any guarantee that $\hat{\mathbf{R}}$ will be positive semi-definite and some sort of remedial measures will be required. For instance, if $\hat{\mathbf{R}}$ is not a positive semi-definite correlation matrix, then

$$\hat{\mathbf{R}}' = \frac{\hat{\mathbf{R}} - \varepsilon \mathbf{I}}{1 - \varepsilon}$$

will be, provided ε is the smallest eigenvalue of $\hat{\mathbf{R}}$ (which will be negative) and \mathbf{I} is the $k \times k$ identity matrix. There are many other fixes; see, for instance, Ghosh and Henderson (2002).

Remark 6.1. Notice that our approach is to decompose the problem into fitting the marginals and the correlations separately, and then solving the correlation-matching problem. Biller and Nelson (2005) describe a least-squares method for fitting the marginals and correlation structure simultaneously.

Generating Multivariate Normal Vectors

A component of the NORTA method is generating standard multivariate normal random vectors with correlation matrix \mathbf{R} . The multivariate normal distribution is completely characterized by its marginal means and standard deviations, and the correlations among all pairs of components. See, for instance, Johnson (1987).

To illustrate how variate generation is done, suppose that $k = 2$ so that we need bivariate normal random vector (Z_1, Z_2) with correlation $r = r_{12}$. If W_1, W_2 are i.i.d. $N(0, 1)$ random variables, then the following transformation works:

$$\begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ r & \sqrt{1-r^2} \end{pmatrix} \begin{pmatrix} W_1 \\ W_2 \end{pmatrix} \quad (6.35)$$

Notice that by definition Z_1 is $N(0, 1)$, while Z_2 is normally distributed because it is a linear combination of normally distributed random variables. Also, $\text{Var}(Z_2) = r^2 + (\sqrt{1-r^2})^2 = 1$. Exercise 37 asks you to prove that $\text{Corr}(Z_1, Z_2) = r$.

Equation (6.35) is based on the fact that the conditional distribution of Z_2 given Z_1 is $N(rZ_1, 1-r^2)$. This insight for $k = 2$ extends to $k > 2$: The conditional distribution of Z_h given Z_1, Z_2, \dots, Z_{h-1} is normal, for $h = 2, 3, \dots, k$. The following algorithm is equivalent to generating Z_1 , then generating Z_2 given Z_1 , then Z_3 given Z_1, Z_2 , etc. (Johnson 1987):

1. Let \mathbf{C} be the Cholesky decomposition of the correlation matrix \mathbf{R} so that $\mathbf{C}^\top \mathbf{C} = \mathbf{R}$.
2. Let $\mathbf{W}^\top = (W_1, W_2, \dots, W_k)$ where the W_i are i.i.d. $N(0, 1)$.
3. Set $\mathbf{Z} = (Z_1, Z_2, \dots, Z_k)^\top = \mathbf{C}^\top \mathbf{W}$.

Remark 6.2. Although not needed to apply the NORTA method, another useful fact is that if $Z \sim N(0, 1)$, then $\mu + \sigma Z \sim N(\mu, \sigma^2)$. This provides a method for generating multivariate normal random vectors with any means, variances and correlations by first generating Z_1, Z_2, \dots, Z_k with the desired correlations, then applying the transformation to achieve the required means and variances.

Solving the Correlation-Matching Problem

The correlation-matching problem is simplified by two insights, one obvious, the other not:

1. For a NORTA vector, the correlation ρ_{ij} is a function of only the base correlation r_{ij} . This is true because

$$\rho_{ij} = \text{Corr}(X_i, X_j) = \text{Corr}\left(F_i^{-1}[\Phi(Z_i)], F_j^{-1}[\Phi(Z_j)]\right)$$

which is only a function of Z_i and Z_j . Thus, the correlation-matching problem decomposes into $k(k-1)/2$ individual matching problems.

2. Under very mild conditions, $\text{Corr}\left(F_i^{-1}[\Phi(Z_i)], F_j^{-1}[\Phi(Z_j)]\right)$ is a continuous, nondecreasing function of r_{ij} (Billar and Nelson 2003). Thus, a straightforward search can be used to find r_{ij} .

To implement the search we need to be able to calculate $\text{Corr}\left(F_i^{-1}[\Phi(Z_i)], F_j^{-1}[\Phi(Z_j)]\right)$ for any given r_{ij} . There are essentially two choices: numerical methods or simulation. We present a simulation approach here because of its simplicity; for numerical methods see Cario and Nelson (1998) and Biller and Nelson (2003).

The simulation method is straightforward: Given a pair of marginals F_i and F_j , a correlation ρ_{ij} to match, and a feasible value of r_{ij} to check, generate m independent replications of $X_i = F_i^{-1}[\Phi(Z_i)]$ and $X_j = F_j^{-1}[\Phi(Z_j)]$, estimate their correlation $\hat{\rho}$, and adjust r_{ij} until $\hat{\rho} \approx \rho_{ij}$. In the algorithm below we implement a bisection search, and we call a function `rho` to compute the correlation estimate $\hat{\rho}$. Notice that the number of replications m and error tolerance $\varepsilon > 0$ must be specified; we comment on these choices later.

Empirical correlation matching

1. generate $\mathbf{W}_1 = (W_{11}, W_{12}, \dots, W_{1m})$ and $\mathbf{W}_2 = (W_{21}, W_{22}, \dots, W_{2m})$, all i.i.d. $N(0, 1)$
2. if $\rho < 0$ then set $\ell = -1$ and $u = 0$
Else set $\ell = 0$ and $u = 1$
3. set $r = \rho_{ij}$
4. set $\hat{\rho} = \text{rho}(r, \mathbf{W}_1, \mathbf{W}_2)$
5. while $|\hat{\rho} - \rho_{ij}| > \varepsilon$ do
 - a. if $\hat{\rho} > \rho_{ij}$ then set $u = r$
Else set $\ell = r$
 - b. set $r = (\ell + u)/2$
 - c. set $\hat{\rho} = \text{rho}(r, \mathbf{W}_1, \mathbf{W}_2)$
6. loop
7. return $r_{ij} = r$

The function below computes the correlation estimate given a value of r :

`rho(r, $\mathbf{W}_1, \mathbf{W}_2$)`

1. set $Z_{1j} = W_{1j}, j = 1, 2, \dots, m$
2. set $Z_{2j} = rZ_{1j} + \sqrt{1 - r^2}W_{2j}, j = 1, 2, \dots, m$
3. set $X_{ij} = F_i^{-1}[\Phi(Z_{ij})], i = 1, 2; j = 1, 2, \dots, m$
4. return

$$\text{rho} = \frac{\sum_{h=1}^m (X_{1h} - \bar{X}_1)(X_{2h} - \bar{X}_2)}{\sqrt{[\sum_{h=1}^m (X_{1h} - \bar{X}_1)^2][\sum_{h=1}^m (X_{2h} - \bar{X}_2)^2]}}$$

Clearly, the choice of sample size m and error tolerance ε matter. We recommend m no less than 10,000 and ε no greater than $0.05|\rho_{ij}|$, with $0.01|\rho_{ij}|$ preferred. Chen (2001) describes how a sequence of experiments with increasing m and decreasing ε can be used to provide controlled error. This is essential because as k gets larger, a poorly estimated \mathbf{R} may not be positive definite, a condition required for it to be a legitimate correlation matrix.

Also notice that only one set of $\mathbf{W}_1, \mathbf{W}_2$ need to be generated.

Evaluation of the CDF and Inverse CDF

The NORTA method as described here requires three things:

1. The capability to evaluate the standard normal cdf, $\Phi(\cdot)$.
2. The capability to evaluate the inverse cdfs of the target marginal distributions, $F_1^{-1}(\cdot), F_2^{-1}(\cdot), \dots, F_k^{-1}(\cdot)$.
3. The capability to generate i.i.d. standard normal random variates W_{ij} .

Many numerical libraries contain efficient methods to evaluate $\Phi(\cdot)$, and computer code can be found in textbooks such as Bratley et al. (1987). The NORTA method is very sensitive to having an accurate evaluation of the normal cdf, and implementations that do not handle the tails in a sensible way can cause the empirical correlation matching to fail to converge.

As discussed in Sect. 6.4, even when $F^{-1}(\cdot)$ cannot be expressed in closed form, it is still possible to numerically invert the cdf. Numerical methods to evaluate the inverse cdf of many common distributions can also be found in numerical libraries as well as Bratley et al. (1987).

Exercises

1. Suppose that Z_1, Z_2, \dots are i.i.d. positive random variables, and $X = \prod_{i=1}^n Z_i$. Let $W = \ln(X) = \sum_{i=1}^n \ln(Z_i)$, and suppose that $E(\ln(Z_1)^2) < \infty$. Use the central limit theorem and the continuous mapping theorem to show that X is asymptotically lognormal.
2. Show that the hazard function for the Weibull distribution is increasing for $\alpha > 1$, decreasing for $\alpha < 1$, and constant for $\alpha = 1$.
3. Derive the relationship between the mean and variance of Y and Y' when $Y = a + (b - a)Y'$.
4. Show that X and X' in (6.9) have the same skewness and kurtosis.
5. What problem could occur if instead of solving the least-squares formulation in (6.10), we instead tried to solve

$$\hat{\theta}_U = \operatorname{argmin}_{\theta} \sum_{i=1}^m w(i) \left(F(X_{(i)}; \theta) - \frac{i}{m} \right)^2 ?$$

Hint: Suppose F is a distribution with unbounded support, such as the exponential distribution.

6. Suppose we want to parameterize a triangular distribution by specifying the minimum, maximum and mean value. Derive the most likely value b given this information.
7. Prove that the ecdf is asymptotically consistent $\widehat{F}(x) \xrightarrow{a.s.} F(x)$.
8. Suppose that $\tilde{X} \sim \tilde{F}$, the ecdf of X_1, X_2, \dots, X_m . Show that

$$\text{Var}(\tilde{X} | X_1, \dots, X_m) = \frac{1}{m} \sum_{i=1}^m (X_i - \bar{X})^2.$$

9. For the linearly interpolated ecdf, show that $E(\tilde{F}(x)) \neq F_X(x)$ in general (a single example will suffice).
10. Suppose that $\tilde{X} \sim \tilde{F}$, the linearly interpolated ecdf. Derive a general expression for $E(\tilde{X} | X_1, \dots, X_m)$ to show that it is not equal to \bar{X} , except as $m \rightarrow \infty$.
11. Suggest ways to model the arrival of fans to a basketball game, or more generally arrivals to an event with a fixed starting time and capacity.
12. In this chapter we focused on nonparametric estimators of $\lambda(t)$ or $\Lambda(t)$. However, there is a substantial literature on parametric models, such as

$$\lambda(t) = \exp \left\{ \sum_{i=0}^p \beta_i t^i + \eta \sin(\omega t + \phi) \right\}$$

(Lee et al. 1991), where p, β_i, η, ω and ϕ are parameters to be estimated from arrival data. This model allows cyclic behavior via $\eta \sin(\omega t + \phi)$ and a long-term trend through $\sum_{i=0}^p \beta_i t^i$. What is the advantage of using this form, rather than the apparently simpler $\lambda(t) = \sum_{i=0}^p \beta_i t^i + \eta \sin(\omega t + \phi)$?

13. Prove that $\tilde{\Lambda}(t) \xrightarrow{a.s.} \Lambda(t)$ and that $E(\tilde{\Lambda}(t)) = \Lambda(t)$ for any fixed $t \in [0, T)$. *Hint:* For any fixed t , $N_1(t), N_2(t), \dots, N_k(t)$ are i.i.d. random variables.
14. For the arrival-rate function $\tilde{\lambda}(t)$ obtained by counting arrivals in fixed intervals of size δ , derive the corresponding integrated rate function.
15. Data on the arrival of faxes, by hour, for 1 month (31 days) can be found on the book website in `FaxCounts.xls`. Use these data to estimate arrival rates for a piecewise-constant, nonstationary arrival process.
16. For the call center in Exercise 16 of Chap. 4, the arrival rate is not actually a steady 60/h, it varies throughout the day. Data on call counts, by hour, for 1 month (31 days) can be found on the book website in `CallCounts.xls`. Use these data to estimate arrival rates for a piecewise-constant, nonstationary Poisson arrival process. Implement this arrival process in your call center simulations. Does this affect your recommendation?
17. Consider the call arrival data from the previous exercise. Let $N(t)$ represent the cumulative number of arrivals by time t . If the process is nonstationary Poisson then $\text{Var}(N(t)) / E(N(t)) = 1$ for all t , or stated differently $\text{Var}(N(t)) = \beta E(N(t))$ with $\beta = 1$. Since you have arrival count data, you can estimate $\text{Var}(N(t))$ and $E(N(t))$ at $t = 1, 2, \dots, 8$ h. Use these data to fit the regression model $\text{Var}(N(t_i)) = \beta E(N(t_i))$ and see if the estimated value of β supports the choice of a nonstationary Poisson arrival process. *Hints:* This

is regression through the origin. Also, remember that $N(t_i)$ represents the *total number of arrivals by time t_i* .

18. Each summer a car dealer runs a 40-h sale. Customer arrival times (in hours) from two of these sales can be found on the book website in `ArrivalTimes.xls`. Use these data to form a linearly interpolated integrated rate function $\hat{\Lambda}(t)$, then generate customer arrival times for 1 day from a nonstationary Poisson arrival process using $\hat{\Lambda}(t)$.
19. For the Weibull distribution with cdf (6.4), derive the inverse cdf in a form useful for random-variate generation.
20. Derive the cdf of the triangular distribution (6.15), then derive the inverse cdf in a form useful for random-variate generation.
21. Show that the inverse cdf for the linearly interpolated cdf is

$$\tilde{X} = X_{(i)} + (m - 1)(X_{(i+1)} - X_{(i)}) \left(U - \frac{i - 1}{m - 1} \right)$$

where $i = \lceil (m - 1)U \rceil$.

22. A simple model of a doctor’s office is a single server queue. Suppose that patient arrivals to the office are schedule at 15-min intervals, but patients deviate from their scheduled time by an amount that can be modeled as normally distributed with mean -5 min and standard deviation 1 min. How can you generate arrivals to such a queue?
23. We claim the following algorithm implements the inverse cdf method to generate Poisson random variates with mean λ ; that is $\Pr\{X = x\} = e^{-\lambda} \lambda^x / x!$ for $x = 0, 1, 2, \dots$

- a. Set $c = d = e^{-\lambda}$ and $X = 0$
- b. Generate $U \sim U(0, 1)$
- c. Until $U \leq c$ do
 - i. $X = X + 1$
 - ii. $d = d\lambda/x$
 - iii. $c = c + d$
- Loop
- d. Return X

First prove that this algorithm does as claimed. Then derive the expected number of times that the loop is executed for each random variate generated. Notice that it is an increasing function of λ , so the algorithm becomes less and less efficient as λ increases. *Hint:* $\Pr\{X = x\} = \lambda \Pr\{X = x - 1\} / x$ for $x \geq 1$.

24. How well does the method of Eq. (6.29) approximate a normally distributed random variate?
25. Prove that $\lambda \int_0^t (1 - G(s)) ds = G(t)$ if $G(t) = 1 - e^{-\lambda t}$. Then show that $G_e \neq G$ in general.

26. Show that if we apply the inversion method to $\bar{\Lambda}(t)$ that (a) only the observed arrival times $T_{(i)}, i = 0, 1, \dots, C$ can be generated, and (b) that we are likely to generate multiple arrivals with the same arrival time.
27. A disadvantage of using $\hat{\Lambda}(t)$ is the need to store all of the arrival times $T_{(i)}, i = 1, 2, \dots, C$. Suppose instead that we only stored $\bar{\Lambda}(\delta), \bar{\Lambda}(2\delta), \dots, \bar{\Lambda}(m\delta)$, where $\delta = T/m$. Propose an interpolated estimator of $\Lambda(t)$ based on these values, and derive a variate-generation algorithm for it.
28. For the inversion algorithm applied to $\hat{\Lambda}(t)$, show that $\tilde{S}_n = C/k$ maps into an arrival at time T .
29. Suppose that arrivals occur according to a NSPP process with integrable rate function $Z\lambda(t)$, conditional on a positive random variable Z with $E(Z) = 1$. In other words, to simulate this process we first generate a value of Z , and then generate arrivals from a NSPP with rate $Z\lambda(t)$ using inversion or thinning. Derive $E(N(t))$ and $\text{Var}(N(t))/E(N(t))$ for this process, and compare it to results for inverting a renewal process. This arrival process is called a *doubly stochastic Poisson process* and it is popular in modeling call center arrivals.
30. Let X be a discrete-valued random variable that takes values $x_0 < x_1 < \dots < x_n$ with probabilities p_0, p_1, \dots, p_n , respectively. Derive, and prove the correctness of a rejection algorithm that first generates an index $I \in \{0, 1, \dots, n\}$ using the discrete uniform distribution. What is the expected number of trials for your algorithm to generate one X ? *Hint*: The following simple algorithm works but is very inefficient: (i) Generate $W \sim U(0, 1)$ and set $I = \lceil (n+1)W \rceil - 1$. (ii) Generate $U \sim U(0, 1)$. (iii) If $U \leq p_I$ then return $X = x_I$; otherwise go to Step (iii). First prove that this algorithm works, and that it has expected number of trials $n+1$. Then derive a better algorithm by applying the concept of a majorizing function as used in the continuous case.
31. Apply your answer to Exercise 30 to produce a rejection algorithm for the binomial distribution on n trials with probability of success $0 < p < 1$. That is,

$$p_X(i) = \Pr\{X = i\} = \binom{n}{i} p^i (1-p)^{n-i}, \quad i = 0, 1, 2, \dots, n.$$

Hint: The mode of the binomial distribution is either $i^* = \lfloor (n+1)p \rfloor$ or $i^* = \lfloor (n+1)p \rfloor - 1$.

32. In the rejection method, why is it necessary to resample both U and Y on each trial? Is the following algorithm sufficient?
- Generate $U \sim U(0, 1)$
 - Generate $V \sim g$
 - If $U \leq f_X(V)/m(V)$, then return $X = V$; else go to Step (b).
33. Consider generating a NSPP with rate $\lambda(t)$ using the following algorithm (with $S_0 = 0$):

$$S_{n+1} = S_n - \ln(1 - U_{n+1})/\lambda(S_n)$$

where U_1, U_2, \dots are i.i.d. $U(0, 1)$. That is, we generate an exponential interarrival time using the arrival rate in force at the most recent arrival time. Prove that this method does not guarantee an arrival process with rate $\lambda(t)$. *Hint:* Consider an arrival rate $\lambda(t)$ that is 0 for some interval of time.

34. The inversion method for generating a nonstationary arrival process requires an equilibrium base process with rate 1 and variance σ_A^2 . For $\sigma_A^2 > 1$, Gerhardt and Nelson (2009) suggest a balanced hyperexponential distribution; this means that \tilde{A} is exponentially distributed with rate λ_1 with probability p , and exponentially distributed with rate λ_2 with probability $1 - p$. “Balance” means that $p/\lambda_1 = (1 - p)/\lambda_2$. Thus, there are only two free parameters, p and λ_1 . Show that we achieve the desired arrival rate and variance if

$$p = \frac{1}{2} \left(1 + \sqrt{\frac{\sigma_A^2 - 1}{\sigma_A^2 + 1}} \right) \quad \lambda_1 = 2p.$$

35. The inversion method for generating a nonstationary arrival process requires an equilibrium base process with rate 1 and variance σ_A^2 . When $\sigma_A^2 < 1$, Gerhardt and Nelson (2009) suggest a mixture of Erlangs of common order. First find an integer k such that $1/k \leq \sigma_A^2 < 1/(k - 1)$. Then with probability p , \tilde{A} is Erlang distributed with $k - 1$ phases and mean $(k - 1)/\lambda$, and with probability $1 - p$ it is Erlang distributed with k phases and mean k/λ . Thus, there are two free parameters, p and λ . Show that we achieve the desired arrival rate and variance if

$$p = \frac{1}{1 + \sigma_A^2} \left(k\sigma_A^2 - \sqrt{k(1 + \sigma_A^2) - k^2\sigma_A^2} \right)$$

$$\lambda = k - p.$$

36. Suppose Y is a continuous-valued random variable with strictly increasing cdf F . Prove that $U = F(Y)$ has the $U(0, 1)$ distribution. This is known as the *probability–integral transformation*.
37. Show that the transformation (6.35) implies that $\text{Corr}(Z_1, Z_2) = r$. *Hint:* Since we have already shown that $\text{Var}(Z_1) = \text{Var}(Z_2) = 1$, then $\text{Corr}(Z_1, Z_2) = \text{Cov}(Z_1, Z_2)$.
38. As part of the assembly process for laptop computers, batches of computers are loaded into a rack and software is copied to their hard drives. The time to load computers into the rack depends on how many computers are in the batch, but is not perfectly predictable. For a simulation of laptop assembly, you need an input model that generates a load time given a batch size. Develop an appropriate input modeling and variate-generation algorithm, carefully justifying your approach. You have the following data:

Batch size	Load time (min)
5	9.9
10	15.2
15	20.2
20	24.4
10	14.4
10	14.8
15	20.8
5	10.0
5	9.0
15	19.5
5	9.9
5	9.4
10	15.8
15	20.1
20	25.7
5	10.5
10	15.6
10	15.2
5	10.2
20	24.5

39. For a random variable X that has density function

$$f(x) = \begin{cases} 0, & x < -1 \\ \frac{x+1}{2}, & -1 \leq x \leq 1 \\ 0, & x > 1 \end{cases}$$

provide two variate-generation algorithms, one using the inverse cdf and the other using rejection.

40. Given observations of arrival times on $[0, T]$ the linearly interpolated integrated rate function is

$$\widehat{\Lambda}(t) = \left(\frac{C}{C+1} \right) \left[\frac{i}{k} + \frac{1}{k} \left(\frac{t - T_{(i)}}{T_{(i+1)} - T_{(i)}} \right) \right], \text{ when } T_{(i)} < t \leq T_{(i+1)}$$

with $T_{(0)} = 0$, $T_{(C+1)} = T$, and $T_{(i)}$ the sorted arrival times. Suppose that we wanted to generate arrivals with this integrated rate function, but by thinning rather than inversion. Thus, we need an arrival rate function $\widehat{\lambda}(t)$. First give an expression for the rate function that is implied by $\widehat{\Lambda}(t)$. Then give an expression for the maximum arrival rate $\tilde{\lambda}$ to use for thinning.

41. Implement the random vector example from Sect. 6.3.2 and compare your results to those in the book which used $m = 10,000$ and $\varepsilon = 0.01|\rho_{12}|$. Try fitting various marginal distributions for which you have access to the inverse cdf (Excel, Matlab, R and other software applications have the inverse cdf of many distributions as built-in functions, along with the cdf of the normal distribution).

- 42. A popular full-period MCG is $z_i = 16,807z_{i-1} \bmod 2^{31} - 1$. Implement this generator, using double-precision floating point arithmetic to avoid overflow problems.
- 43. Starting with $z_0 = 1$, generate ten additional starting seeds 100,000 values apart for the generator in Exercise 42 above.
- 44. Show that a MRG of order K can be represented as

$$\mathbf{z}_i = \mathbf{A}\mathbf{z}_{i-1} \bmod m$$

where $\mathbf{z}_i = (z_{i-K+1}, z_{i-K}, \dots, z_i)^\top$ and

$$\mathbf{A} = \begin{pmatrix} 0 & 1 & \dots & 0 \\ \vdots & & \ddots & 0 \\ 0 & 0 & \dots & 1 \\ a_k & a_{k-1} & \dots & a_1 \end{pmatrix}.$$

Also show that we can skip ahead q values by using the relationship

$$\mathbf{z}_{i+q} = \mathbf{A}^q \mathbf{z}_{i-1} \bmod m = (\mathbf{A}^q \bmod m) \mathbf{z}_{i-1} \bmod m.$$

- 45. For a MCG show that $z_i = a^i z_0 \bmod m$. *Hint:* If b_1, b_2, m are integers, and $r_i = b_i \bmod m$, then $(r_1 r_2 \bmod m) = (b_1 b_2 \bmod m)$.
- 46. For the following combined linear congruential generator of the L'Ecuyer type

$$Z_i = (Z_{i,1} - Z_{i,2}) \bmod 15$$

$$U_i = \begin{cases} \frac{Z_i}{16}, & Z_i > 0 \\ \frac{15}{16}, & Z_i = 0 \end{cases}$$

where the two generators are

$$\begin{aligned} Z_{i,1} &= (11Z_{i-1,1}) \bmod 16 && \text{with } Z_{0,1} = 1 \\ Z_{i,2} &= (2Z_{i-1,2}) \bmod 13 && \text{with } Z_{0,2} = 1 \end{aligned}$$

generate U_1 and U_2 by hand. Next implement the generator.

Chapter 7

Simulation Output

This chapter describes simulation output analysis, specifically estimating means, probabilities and quantiles of simulation output, along with measures of error on these estimates. The distinction between risk and error is emphasized, along with the impact of uncertainty about the input processes that drive the simulation.

7.1 Performance Measures, Risk and Error

Figure 7.1 displays a histogram and empirical cdf (ecdf) of 1,000 replications of the time to complete the project planning stochastic activity network (SAN) described in Sect. 3.4 and simulated in Sect. 4.4. Let the random variable Y denote the time to complete the project with cdf F_Y ; $Y_1, Y_2, \dots, Y_{1,000}$ are the results from 1,000 i.i.d. replications; $Y_{(1)} \leq Y_{(2)} \leq \dots \leq Y_{(1,000)}$ are the sorted values (called *order statistics*); and \hat{F} is the ecdf. From these figures we can visualize estimators of several system performance measures:

- The expected value of the time to complete the project (the mean), $\mu = E(Y)$, is estimated by the sample mean $\bar{Y} = \sum_{i=1}^n Y_i/n$.
- The probability that the project completes in no more than 5 days, $\theta = F_Y(5)$, is estimated by $\hat{F}(5) = \#\{Y_i \leq 5\}/1,000$.
- The 0.95 quantile of the time to complete the project, $\vartheta = F_Y^{-1}(0.95)$, is estimated by $\hat{F}^{-1}(0.95) = Y_{(950)}$. The 0.95 quantile is a time ϑ such that the probability of completing the project by or before that time is 0.95.

This section covers estimating and interpreting these performance measures based on simulation output.

For the SAN, the probability θ and the quantile ϑ are more relevant performance measures than the mean μ . For whoever requested the project, θ is the probability that it completes by a desired date. For whoever is executing the project, ϑ is a completion date that they can promise with a given probability of being able to meet it. The mean time to complete the project is less informative: Assuming that

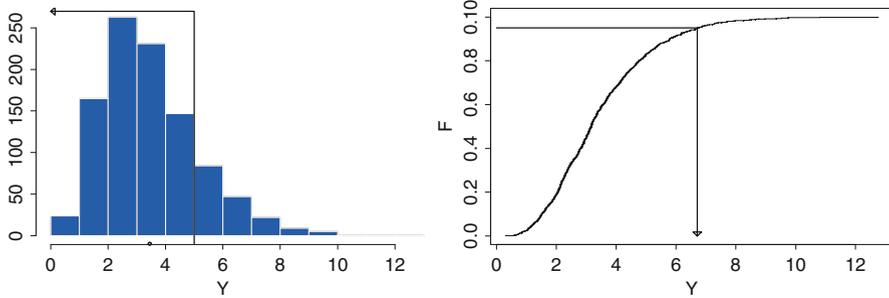


Fig. 7.1 Histogram (*left*) and ecdf (*right*) of 1,000 replications of the time to complete the stochastic activity network. Histogram: Project completion times ≤ 5 are indicated by the *arrow*; the sample mean by the *circle*. Ecdf: $Y(950)$ is indicated by the *arrow*

the project will only be undertaken once, the histogram indicates that a completion time near the mean is possible but not at all certain.

By way of contrast, for the hospital reception described in Sect. 3.2 and simulated in Sect. 4.3, the focus was on μ , the long-run average (mean) patient or visitor waiting time. The mean is clearly relevant here as it summarizes performance averaged over the very large number of patients and visitors who will arrive over time. Less obvious is what measures are most relevant for the Asian option described in Sect. 3.5 and simulated in Sect. 4.5. While it may seem similar to the SAN example, since we will likely only buy or sell this option once, financial theory shows that under certain assumptions about the market the price of an option should be its expected payoff under a particular probability model; see, for instance, Glasserman (2004, Sect. 1.2). Therefore, valuing the option by its expected value v makes sense.

A point of emphasis in this chapter is that some performance measures, such as the mean μ , are most useful in characterizing a system over the long run, where “long run” implies a large number of occurrences. Other measures, such as a probability θ or quantile ϑ can be used as measures of *risk*, helping to quantify what is likely to happen “the next time” rather than “over the long run.” Of course, μ , θ and ϑ all have to be estimated, and estimates are subject to sampling error. Measures of sampling error are another theme of this chapter; they indicate whether we have done enough simulation (e.g., obtained enough replications) to trust the performance estimates we have. This chapter focuses on measures of error for estimators that are functions of i.i.d. data from across replications; results for dependent data from a single replication of a steady-state simulation are addressed in Chap. 8.

7.1.1 Point Estimators and Measures of Error

Point estimators that are sample means, or functions of sample means, were analyzed in Sect. 5.2. By way of review, and keeping the SAN data in Fig. 7.1 in mind,

if we want to estimate $\mu = E(Y)$ from i.i.d. observations with finite variance σ^2 , then the natural unbiased estimator is the sample mean \bar{Y} which has standard error $se(\bar{Y}) = \sigma/\sqrt{n}$. For large n the central limit theorem justifies the $(1 - \alpha)100\%$ confidence interval (CI) for μ

$$\bar{Y} \pm z_{1-\alpha/2} \frac{S}{\sqrt{n}} \quad (7.1)$$

where S is the sample standard deviation, computed as the square root of the sample variance

$$S^2 = \frac{1}{n-1} \sum_{i=1}^n (Y_i - \bar{Y})^2 \quad (7.2)$$

$$= \frac{1}{n-1} \left[\sum_{i=1}^n Y_i^2 - \frac{1}{n} \left(\sum_{j=1}^n Y_j \right)^2 \right]. \quad (7.3)$$

The expression in Eq. (7.3) is the one that is most often implemented because it allows S^2 can be computed in one pass through the data, while (7.2) requires two passes.

For the data in Fig. 7.1 a 95% CI for μ is 3.46 ± 0.11 . The ± 0.11 is a measure of error for our point estimate of the mean time to complete the project, 3.46 days. If we increased the number of replications we would expect the width of the CI to decrease as $1/\sqrt{n}$.

Estimates of probabilities such as $\theta = \Pr\{Y \leq y\}$ can also be viewed as sample means, since the ecdf evaluated at any fixed value of y is the average of indicator functions:

$$\hat{F}(y) = \frac{1}{n} \sum_{i=1}^n I(Y_i \leq y). \quad (7.4)$$

Thus, for large n the CI (7.1) still applies and careful algebra shows that

$$S^2 = \left(\frac{n}{n-1} \right) \hat{F}(y) (1 - \hat{F}(y)). \quad (7.5)$$

When n is large the ratio $n/(n-1)$ is often treated as 1.

For the data in Fig. 7.1 a 95% CI for θ is 0.17 ± 0.02 . Again, the ± 0.02 is a measure of error for our point estimate of the probability that the project completes in 5 days or less, 0.17. In this case we actually know that $\theta = 0.165329707$, well within the measure of error.

The standard error of a probability estimator $\hat{\theta}$ is

$$se(\hat{\theta}) = \sqrt{\frac{\theta(1-\theta)}{n}}$$

which is maximized at $\theta = 1/2$. However, for probabilities relative error is more revealing:

$$\frac{se(\hat{\theta})}{\theta} = \sqrt{\frac{1-\theta}{n\theta}}.$$

As $\theta \rightarrow 0$ the size of the estimation error relative to the probability that we are estimating increases dramatically; in other words, the more rare the event is, the more difficult it is to estimate its probability precisely.

Remark 7.1. In the discussion above we assumed that $\theta \neq 0$. See Lewis (1981) for the case when the event of interest is never observed to occur on any observation of the simulation experiment.

Suppose that the cdf of Y , F_Y , is strictly increasing and has a density f_Y . The q -quantile is the value ϑ such that $F_Y(\vartheta) = q$, or equivalently $\vartheta = F_Y^{-1}(q)$. Clearly the natural estimator $\hat{\vartheta} = \hat{F}^{-1}(q) = Y_{(\lceil nq \rceil)}$ is not a sample average, so its distributional properties are more complicated. Nevertheless it can be shown that $\hat{\vartheta}$ does satisfy a central limit theorem:

$$\sqrt{n}(\hat{\vartheta} - \vartheta) \xrightarrow{D} N\left(0, \frac{q(1-q)}{[f_Y(\vartheta)]^2}\right) \quad (7.6)$$

as $n \rightarrow \infty$ provided $f_Y(\vartheta) > 0$.¹ Thus, a large n approximation of the standard error is

$$\text{se}(\hat{\vartheta}) \approx \sqrt{\frac{q(1-q)}{n[f_Y(\vartheta)]^2}}.$$

Unfortunately, even though the standard error of $\hat{\vartheta}$ decreases as $1/\sqrt{n}$, such as that of a sample average, this result is not easily applied since it involves the unknown density function f_Y evaluated at the unknown quantile ϑ . However, this result is helpful for seeing the impact of q on the standard error, as shown in the following example.

Suppose $f_Y(y) = e^{-y}$, $y \geq 0$, the exponential distribution with mean 1. Then $\vartheta = -\ln(1-q)$, so that

$$\text{se}(\hat{\vartheta}) \approx \sqrt{\frac{q(1-q)}{n[\exp(\ln(1-q))]^2}} = \sqrt{\frac{q}{n(1-q)}}.$$

Thus, the standard error of $\hat{\vartheta} = Y_{(\lceil nq \rceil)}$ as an estimator of the q quantile increases dramatically as q approaches 1, meaning a quantile farther out in the right tail. For instance, when Y is exponentially distributed the standard error of $\hat{\vartheta}$ for estimating the 0.99 quantile is roughly ten times larger than the standard error for estimating the 0.5 quantile (the median), for the same n . Therefore, the more extreme the quantile is the more difficult it is to estimate.

Fortunately, there is a CI for quantiles that is not based on the central limit theorem. Notice that since $\Pr\{Y_i \leq \vartheta\} = q$, and the observations Y_i are i.i.d., then

¹ This central limit theorem is obtained by noting that $\Pr\{\sqrt{n}(\hat{\vartheta} - \vartheta) \leq y\} = \Pr\{\hat{\vartheta} \leq \vartheta + y/\sqrt{n}\} = \Pr\{\hat{F}(\vartheta + y/\sqrt{n}) > q\}$ and then using the fact that the central limit theorem for averages applies to \hat{F} .

$$\#\{Y_i \leq \vartheta\} \sim \text{Binomial}(n, q).$$

This implies that if we find integers $0 \leq \ell < u \leq n$ such that

$$\Pr\{Y_{(\ell)} \leq \vartheta < Y_{(u)}\} = \sum_{i=\ell}^{u-1} \binom{n}{i} q^i (1-q)^{n-i} \approx 1 - \alpha$$

then $[Y_{(\ell)}, Y_{(u)}]$ forms an approximate $(1 - \alpha)100\%$ CI for ϑ . In words, we find (ℓ, u) such that the probability that at least $n - u + 1$ of the largest Y_i 's are greater than ϑ , and at least ℓ of the smallest Y_i 's are less than or equal to ϑ , is approximately $1 - \alpha$.

When n is large and neither nq nor $n(1 - q)$ is small, the normal approximation to the binomial distribution gives easy approximations for ℓ and u :

$$\begin{aligned} \hat{\ell} &= \left\lfloor nq - z_{1-\alpha/2} \sqrt{nq(1-q)} \right\rfloor \\ \hat{u} &= \left\lceil nq + z_{1-\alpha/2} \sqrt{nq(1-q)} \right\rceil. \end{aligned} \quad (7.7)$$

For instance, with $q = 0.95$, $n = 1,000$ and $z_{0.975} = 1.96$, an approximate 95% confidence interval for the 0.95 quantile is given by $[Y_{(936)}, Y_{(964)}]$. Notice that the values of $\ell = 936$ and $u = 964$ are independent of the actual outputs. For the data in Fig. 7.1, the point estimate is 6.71 days with 95% CI [6.43, 7.05].

7.1.2 Measures of Risk and Error

The distinction between risk and error in simulation output analysis is important, and often misunderstood.

- *Measures of risk directly support decision making.* If $F_Y(5)$, which is the probability of completing the SAN project in 5 days or less, is too small, then we might decide not to undertake the project, or to commit additional resources to insure its completion. Similarly, if the 0.95 quantile of completion time, $F_Y^{-1}(0.95)$, is too many days then we might not be willing to bid on the project. The standard deviation of project completion time, σ , is also a measure of risk since it quantifies the average deviation of the actual project completion time from its mean.
- *Measures of error directly support experiment design.* They tell us if we have expended enough simulation effort (e.g., replications) to be confident in our estimates of system performance.

Figure 7.2 displays a MORE (measure of risk and error) plot of the first 100 replications of the SAN output from Fig. 7.1. The MORE plot adds to the histogram a central arrow indicating the sample average, and two other arrows indicating the 0.05 and 0.95 quantile estimates (other quantiles could be used). Thus, the box over the histogram contains 90% of the simulated outcomes that are most central, and so it is labeled as ‘‘Likely.’’ The areas beyond the box are labeled ‘‘Unlikely’’ indicating that they are possible, but more extreme.

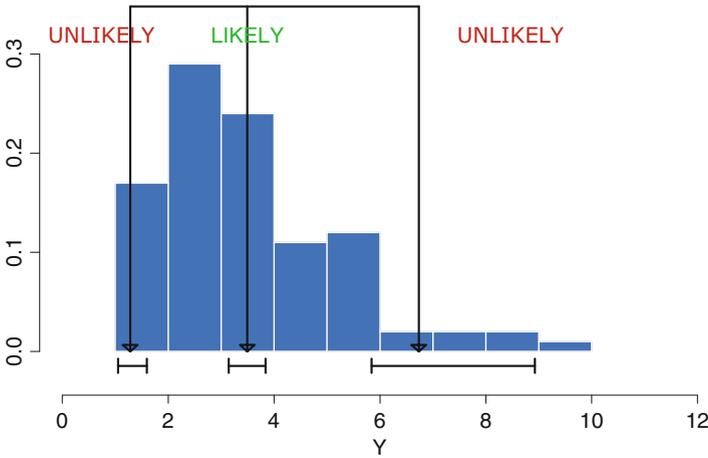


Fig. 7.2 MORE plot of first 100 replications of time to complete the SAN project

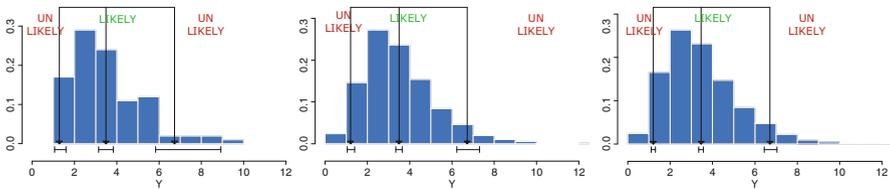


Fig. 7.3 MORE plots for the SAN simulation with (from left to right) $n = 100, 500, 1,000$ replications

Of course, the sample mean and sample quantiles are only estimates of the true mean and quantiles, so the MORE plot also includes confidence intervals for each of these displayed as intervals just below the arrowheads. For instance, Fig. 7.2 shows that the true 0.95 quantile could have substantial error and we should be cautious about making decisions based on it.

Figure 7.3 shows a sequence of MORE plots of the SAN data moving from $n = 100$ to $n = 500$ to $n = 1,000$ replications. Notice that the error in the estimates of the mean and the quantiles decreases, but risk remains, which illustrates that we can simulate away error, but not risk. This plot is useful for displaying both risk and error in an intuitive way. See Nelson (2008) for additional details on the MORE plot.

7.2 Input Uncertainty and Output Analysis

The output analysis methods described in the previous section implicitly assumed that the performance parameters of the simulation model (e.g., the 0.95 quantile

of the simulated time to complete the project) are also the performance parameters of the real-world system in which we are interested (e.g., the real-world project); therefore they only measure estimation error.

Section 5.1 described other sources of error that arise in simulation modeling and analysis, including *input uncertainty*, which refers to the unknown difference between the input models used in the simulation and those that best characterize the real-world system. We focus on the case when the input models have been derived from real-world data, as described in Sects. 6.1–6.3. We first illustrate that input uncertainty can be significant, and then provide one way to capture it.

7.2.1 Input Uncertainty: What Is It?

Consider an $M/M/\infty$ queue, which is a special case of the parking lot problem described in Sect. 3.1 when the arrival rate does not change over time, but instead is a constant λ customers per time. We again denote the mean service time by τ . Much is known about this queue, in particular if Y denotes the steady-state number of customers in the system, then Y has a Poisson distribution with mean $\lambda \tau$.

Suppose λ and τ are not known, so we observe m i.i.d. interarrival times, A_1, A_2, \dots, A_m , and m i.i.d. service times X_1, X_2, \dots, X_m from the “real world” and use them to fit input models. Specifically, we estimate λ by

$$\hat{\lambda} = \left(\frac{1}{m} \sum_{i=1}^m A_i \right)^{-1}$$

and τ by

$$\hat{\tau} = \frac{1}{m} \sum_{i=1}^m X_i.$$

That is, we estimate the arrival rate as the inverse of the sample mean time between arrivals, and the mean service time by the sample mean service time. We assume that we are confident that the distributions of A and X are exponential, but are uncertain about the values of the parameters.

To allow for an analysis, suppose that when we simulate this queue we record a single observation of Y (the number in the queue) in steady state on each of n replications, say Y_1, Y_2, \dots, Y_n (typically we would observe $Y(t)$ for some period of time and take a time average). We then estimate the steady-state mean number in queue by the sample mean

$$\bar{Y} = \frac{1}{n} \sum_{i=1}^n Y_i.$$

Clearly $E(\bar{Y} | \hat{\lambda}, \hat{\tau}) = \hat{\lambda} \hat{\tau}$ and $\text{Var}(\bar{Y} | \hat{\lambda}, \hat{\tau}) = \hat{\lambda} \hat{\tau} / n$. Why? Because conditional on $\hat{\lambda}$ and $\hat{\tau}$ each observation Y is Poisson with mean $\hat{\lambda} \hat{\tau}$, and the mean equals the variance for a Poisson random variable. However, since the input parameters are estimates,

we can be almost sure that $\widehat{\lambda} \neq \lambda$ and $\widehat{\tau} \neq \tau$; that is, our parameter estimates are not the true real-world values. What is the impact of this input uncertainty?

In the appendix to this chapter we show that

$$E(\bar{Y}) = \frac{m}{m-1} \lambda \tau \quad (7.8)$$

$$\begin{aligned} \text{Var}(\bar{Y}) &= \frac{m}{n(m-1)} \lambda \tau + \frac{m(2m-1)(\lambda \tau)^2}{(m-1)^2(m-2)} \\ &\approx \frac{\lambda \tau}{n} + \frac{2(\lambda \tau)^2}{m} \end{aligned} \quad (7.9)$$

where the mean and variance are with respect to both the distributions of $\widehat{\lambda}$ and $\widehat{\tau}$ (the input uncertainty) and the Poisson distribution of the simulation output (the estimation error). From Eqs. (7.8)–(7.9) we see that \bar{Y} is biased. Also, its variance has two components, a familiar one that depends on the number of simulation replications n , and an unfamiliar one that depends on the amount of real-world data m used to “fit” the input processes. In this example the bias is slight even with a modest sample of input data (about 1% if we have $m = 100$ real-world observations). However, the variance due to input uncertainty (the second term in (7.9)) could easily dominate the simulation variance; worse, this source of variance is not mathematically derivable for realistic problems. In the next section we provide a simple method to obtain a rough estimate of the variance due to input uncertainty.

7.2.2 Input Uncertainty: What to Do

Estimation error is something we control through simulation effort (e.g., the number of replications, n). Input uncertainty, on the other hand, depends on the quantity of real-world data we have and it may be impossible or impractical to obtain more. So our goal will be to estimate how large input uncertainty is relative to estimation error; the larger it is the more care we should take in using the simulation results. In terms of the example of the previous section, and in particular the variance decomposition in Eq. (7.9), we would like to know if the term $2(\lambda \tau)^2/m$, which corresponds to input uncertainty, is large relative to $\lambda \tau/n$, which represents simulation error.

To simplify the explanation, suppose for the moment that our simulation model has a single input distribution, F_X , from which we have an i.i.d. sample of real-world data X_1, X_2, \dots, X_m . The true input distribution F_X will be estimated by \widehat{F} , a function of this data, which might be the empirical distribution or a parametric distribution with unknown parameters as in the $M/M/\infty$ example of the previous section.

We now represent the output of the simulation on replication j , implemented using estimated input distribution \widehat{F} , as

$$Y_j = \mu(\widehat{F}) + \varepsilon_j \quad (7.10)$$

where $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n$ are i.i.d. with mean 0 and variance σ_S^2 representing the natural variability from replication to replication in the simulation. The mean term, $\mu(\widehat{F})$, depends on what input model we actually used in the simulation.

For instance, in the $M/M/\infty$ example we simulated the service times as exponential with mean $\widehat{\tau}$ when what we really wanted was exponential with mean τ ; and the arrival process as Poisson with rate $\widehat{\lambda}$ when it should be λ . So for the example $\mu(\widehat{F}) = \widehat{\lambda}\widehat{\tau}$, which depends on the random sample of real-world data.

Suppose that instead of one sample of real-world data, we had b independent samples, $X_{i1}, X_{i2}, \dots, X_{im}, i = 1, 2, \dots, b$, and from the i th sample we estimated an input model \widehat{F}_i . We could then imagine simulating n replications using *each* input model, leading to the output

$$Y_{ij} = \mu(\widehat{F}_i) + \varepsilon_{ij} \quad (7.11)$$

for $i = 1, 2, \dots, b$ and $j = 1, 2, \dots, n$.

Model (7.11) is known as a *random-effects model* (see, e.g., Montgomery 2009), because the mean term $\mu(\widehat{F}_i)$ is random, in this case depending on the particular sample $X_{i1}, X_{i2}, \dots, X_{im}$ that we happened to use to estimate F_X . We can characterize input uncertainty by the value of $\sigma_I^2 = \text{Var}[\mu(\widehat{F}_i)]$, which is analogous to the second term in (7.9). A big simplifying approximation we make is that σ_S^2 , the variance of the simulation output, does not depend on which \widehat{F}_i we have; this will rarely be precisely true, but will be acceptable to get an approximation of the effect of input uncertainty.

For Model (7.11) an estimator of σ_I^2 is

$$\widehat{\sigma}_I^2 = \frac{\widehat{\sigma}_T^2 - \widehat{\sigma}_S^2}{n} \quad (7.12)$$

where

$$\widehat{\sigma}_T^2 = \frac{n}{b-1} \sum_{i=1}^b (\bar{Y}_i - \bar{Y}_{..})^2$$

and

$$\widehat{\sigma}_S^2 = \frac{1}{b(n-1)} \sum_{i=1}^b \sum_{j=1}^n (Y_{ij} - \bar{Y}_i)^2$$

(Montgomery 2009). A “ \cdot ” subscript indicates averaging over that index. Intuitively, $\widehat{\sigma}_T^2$ measures both input and simulation variability, so by subtracting out simulation variability we are left with variance due to input uncertainty. In Exercise 17 you are asked to show that $E(\widehat{\sigma}_I^2) = \sigma_I^2$ under Model (7.11). The ratio $\widehat{\sigma}_I^2 / (\widehat{\sigma}_S^2/n)$ is a measure of how large input uncertainty is relative to estimation error.

The practical problem with this proposal is the need for multiple real-world samples. And if we really had this extra data we would almost certainly use all mb observations together to better estimate F_X . However, a statistical technique known as *bootstrapping* allows us to imitate the effect of having multiple real-world samples, even though we only have one (see Efron and Tibshirani (1993) for a general reference on bootstrapping).

Again let $\{X_1, X_2, \dots, X_m\}$ be the one real-world sample of data. To implement the procedure above, we generate b bootstrap samples, each of size m , by sampling m times *with replacement* from $\{X_1, X_2, \dots, X_m\}$. As is typical in the bootstrapping literature, we denote the i th bootstrap sample as $X_{i1}^*, X_{i2}^*, \dots, X_{im}^*$. We then fit an input distribution \widehat{F}_i^* to this bootstrap sample, using the same method we used with the real-world sample, and implement the procedure above by letting \widehat{F}_i^* stand in for \widehat{F}_i . In a problem with more than one input model, such as the $M/M/\infty$ which has two, we do bootstrapping for each distribution separately. Notice that there is no requirement that we have equal numbers of real-world observations for each input model.

Here is the procedure in algorithm form:

1. Given real-world data $\{X_1, X_2, \dots, X_m\}$, do the following:
2. For i from 1 to b
 - a. Generate the bootstrap sample $X_{i1}^*, X_{i2}^*, \dots, X_{im}^*$ by sampling m times with replacement from $\{X_1, X_2, \dots, X_m\}$.
 - b. Fit \widehat{F}_i^* to $X_{i1}^*, X_{i2}^*, \dots, X_{im}^*$. (If there is more than one input model, do Steps 2a and 2b for each one.)
 - c. Simulate n replications $Y_{ij}, j = 1, 2, \dots, n$ using input model(s) \widehat{F}_i^* .
3. Estimate σ_f^2 using (7.12).

The value of b should be at least 10. The number of replications n should be reasonable for the simulation problem if we were not evaluating input uncertainty.

As an example, consider the $M/M/\infty$ queue from the previous section. If $\lambda = 5$ customers/min, $\tau = 1$ min, we observe $m = 100$ real-world interarrival times and $m = 100$ real-world service times, and we make $n = 10$ replications of the simulation, then Eq. (7.9) gives

$$\text{Var}(\bar{Y}) \approx \frac{\lambda\mu}{n} + \frac{2(\lambda\tau)^2}{m} = \frac{5}{10} + \frac{50}{100} \approx \frac{\sigma_S^2}{n} + \sigma_f^2$$

where the last \approx holds if our random-effects model is approximately correct. Running the procedure with $b = 100$ bootstrap samples we obtained $\widehat{\sigma}_S^2 = 5.321$ and $\widehat{\sigma}_f^2 = 0.546$, both very close to the correct values of 5 and 0.5, respectively. Since $\widehat{\sigma}_S^2/10 = 0.5321$, we see that input uncertainty is approximately as large as estimation error. Whether or not this is a concern depends on whether or not a standard error of $\sqrt{0.5321} = 0.729$ is considered large or small for this problem, since the overall error is approximately double this.

Remark 7.2. Because we estimate σ_f^2 by subtracting in Eq. (7.12) it is possible that the estimate will be negative even though variances cannot be negative. A negative $\widehat{\sigma}_f^2$ should be interpreted as implying that the variability due to input uncertainty is small relative to simulation error. See Ankenman and Nelson (2012) for additional discussion and also a method for identifying which input models account for the most input uncertainty.

Appendix: Input Uncertainty for the $M/M/\infty$ Queue

Here we derive Eqs. (7.8) and (7.9). We will use the following results repeatedly: Suppose E_1, E_2, \dots, E_m are i.i.d exponential each with rate ν (mean $1/\nu$), and let $G = \sum_{i=1}^m E_i$. Then G has an Erlang distribution (which is a special case of a gamma distribution), and $G^{-1} = 1/G$ has an inverse gamma distribution. Thus, $m\hat{\tau}$ has an Erlang distribution, while $\hat{\lambda}m$ has an inverse gamma distribution. For these distributions it is known that

$$\begin{aligned} E(G) &= \frac{m}{\nu} \\ \text{Var}(G) &= \frac{m}{\nu^2} \\ E(G^{-1}) &= \frac{\nu}{m-1} \text{ for } m > 1 \\ \text{Var}(G^{-1}) &= \frac{\nu^2}{(m-1)^2(m-2)} \text{ for } m > 2. \end{aligned}$$

These results imply that $E(G^2) = \text{Var}(G) + E^2(G) = (m+1)/(m\nu^2)$ and $E(G^{-2}) = \text{Var}(G^{-1}) + E^2(G^{-1}) = \nu^2/((m-1)(m-2))$. Therefore, the expected value of the point estimator is

$$\begin{aligned} E(\bar{Y}) &= E \left[E \left(\bar{Y} | \hat{\lambda}, \hat{\tau} \right) \right] \\ &= E \left[\hat{\lambda} \hat{\tau} \right] \\ &= E \left[\hat{\lambda} \right] E \left[\hat{\tau} \right] \\ &= E \left[\frac{m}{\sum_{i=1}^m A_i} \right] \times \tau \\ &= \frac{m}{m-1} \lambda \tau. \end{aligned} \tag{7.13}$$

Notice that in Eq. (7.13) the inner expectation averages over the simulation output distribution, while the outer expectation averages over the possible real-world samples. This result shows that the point estimator is, in general, biased when we must estimate the input models.

To derive the variance, we use the variance decomposition

$$\text{Var}(\bar{Y}) = E \left[\text{Var}(\bar{Y} | \hat{\lambda}, \hat{\tau}) \right] + \text{Var} \left[E(\bar{Y} | \hat{\lambda}, \hat{\tau}) \right].$$

Recall that $E(\bar{Y} | \hat{\lambda}, \hat{\tau}) = \hat{\lambda} \hat{\tau}$ and $\text{Var}(\bar{Y} | \hat{\lambda}, \hat{\tau}) = \hat{\lambda} \hat{\tau} / n$. Thus, using our previous result for the mean we get

$$\mathbb{E} \left[\text{Var}(\bar{Y} | \hat{\lambda}, \hat{\tau}) \right] = \mathbb{E} \left[\frac{\hat{\lambda} \hat{\tau}}{n} \right] = \frac{m}{n(m-1)} \lambda \tau.$$

The second term is

$$\begin{aligned} \text{Var} \left[\mathbb{E}(\bar{Y} | \hat{\lambda}, \hat{\tau}) \right] &= \text{Var} \left[\hat{\lambda} \hat{\tau} \right] \\ &= \mathbb{E} \left[(\hat{\lambda} \hat{\tau})^2 \right] - \mathbb{E} \left[\hat{\lambda} \hat{\tau} \right]^2 \\ &= \mathbb{E} \left[\hat{\lambda}^2 \right] \mathbb{E} \left[\hat{\tau}^2 \right] - \left(\frac{m}{m-1} \right)^2 (\lambda \tau)^2 \\ &= \frac{m^2 \lambda^2}{(m-1)(m-2)} \times \frac{(m+1)\tau^2}{m} - \left(\frac{m}{m-1} \right)^2 (\lambda \tau)^2. \end{aligned}$$

The result then follows by collecting terms.

Exercises

1. Show that Eq. (7.5), which gives the sample variance of a probability estimator, is correct.
2. Suppose we observe i.i.d. simulation outputs Y_1, Y_2, \dots, Y_n , each with mean $\mu < \infty$ and variance $\sigma^2 < \infty$, and our goal is to estimate σ^2 using the usual sample variance S^2 given in (7.2). Assuming also that $\mathbb{E}(Y^k) < \infty, k = 3, 4$, show that S^2 satisfies a central limit theorem. *Hint:* This is not an easy problem, so start with the easier problem of proving that

$$\tilde{S}^2 = \frac{1}{n} \sum_{i=1}^n (Y_i - \mu)^2$$

satisfies a central limit theorem, then express S^2 as a function of \tilde{S}^2 . You will need the tools from Sect. 5.2.4.

3. For the ecdf \hat{F} based on n observations Y_1, Y_2, \dots, Y_n , show that $\hat{F}^{-1}(q) = Y_{(\lfloor nq \rfloor)}$ for $0 < q \leq 1$.
4. Derive (and perhaps refine) the normal approximation (7.7) used to form a confidence interval for a quantile.
5. Recall the stochastic activity network simulation of Sect. 4.4. Estimate the 0.95 quantile of the time to complete the project and put a 90% confidence interval on it. Use at least 1,000 replications.
6. Recall the Asian option simulation of Sect. 4.5. Estimate the 0.05 quantile of its value and a 90% confidence interval on it. Use at least 10,000 replications.
7. Recall the Asian option simulation of Sect. 4.5 where the volatility was $\sigma^2 = (0.3)^2$. With all other parameters unchanged, examine the impact of volatility

on the value of the option for $\sigma = 0.1, 0.2, 0.3, 0.4, 0.5$. Estimate the value to within 1% relative error.

- There are many types of options in addition to the Asian option described in Sect. 3.5 and simulated in Sect. 4.5. A *lookback call option* has the payoff

$$X(T) - \min_{0 \leq t \leq T} X(t)$$

which can be thought of as the profit at time T from simultaneously buying the asset at its lowest price during the period 0 to T and selling it for the final price at time T . Thus, its value is

$$E \left[e^{-rT} \left(X(T) - \min_{0 \leq t \leq T} X(t) \right) \right].$$

Using the same asset model as the Asian option with $T = 1$, $X(0) = \$50$, $r = 0.05$ and $\sigma^2 = (0.3)^2$, estimate the value of this option to within 1% relative error. Notice that you will have to use a discrete approximation for $\min_{0 \leq t \leq T} X(t)$ taking steps of size $\Delta t = T/m$. Experiment with $m = 8, 16, 32, 64, 128$ steps to see how it effects the estimated value, and provide intuition for the effect you see.

- There are many types of options in addition to the Asian option described in Sect. 3.5 and simulated in Sect. 4.5. A *down-and-out call option* with barrier B and strike price K has the payoff

$$I \left\{ \min_{0 \leq t \leq T} X(t) > B \right\} (X(T) - K)^+$$

which means that if the value of the asset falls below $\$B$ before the option matures then the option is worthless. Thus, its value is

$$E \left[e^{-rT} I \left\{ \min_{0 \leq t \leq T} X(t) > B \right\} (X(T) - K)^+ \right].$$

Using the same asset model as the Asian option with $T = 1$, $X(0) = \$50$, $K = \$55$, $r = 0.05$ and $\sigma^2 = (0.3)^2$, estimate the value of this option to within 1% relative error. Notice that you will have to use a discrete approximation for $\min_{0 \leq t \leq T} X(t)$ taking steps of size $\Delta t = T/m$. Experiment with $m = 8, 16, 32, 64, 128$ steps and barrier $B = 30, 35, 40, 45$ to see how they effect the estimated value, and provide intuition for the effect you see.

- A small warehouse supplies paper products. Trucks pull up to the loading bays, the driver places an order and then waits for the order to be gathered by a warehouse worker with a forklift truck. There are currently four loading bays, and when they are full, trucks wait in the parking lot next to the warehouse.

The warehouse is open from 7 a.m. to 4 p.m. Trucks arrive throughout the day. The warehouse has two forklift drivers who pick orders one-at-a-time, on a first-come-first-served basis. The truck drivers load their own orders after the

entire order is picked (for safety reasons), and then leave the loading dock. As soon as the forklift driver has picked the order for a truck they are free to start on another truck if there is one waiting in a loading bay (i.e., they do not have to wait for the driver to load the order).

The warehouse company is interested in how much improvement in efficiency they can obtain from three changes: (a) adding another forklift and driver; (b) adding another loading bay; or (c) doing both (a) and (b). The least expensive option is (a), but they might be willing to do (b) or (c) if there is enough value.

“Efficiency” will be measured in terms of the daily mean number of trucks waiting in the parking lot, mean time a truck spends waiting in the parking lot, mean time a truck waits for a forklift after entering a loading bay, utilization of the bays and the forklifts, and the mean number of trucks still in the parking lot at 4 p.m. (trucks that arrive before 4 p.m. will be served, so having a large number of trucks waiting at 4 p.m. implies that the company will have to pay a great deal of overtime).

Trucks arrive according to a Poisson arrival process with mean time between arrivals of 18 min. The time to pick orders by the forklift drivers is modeled as Erlang with mean 40 min and four phases; the time for a driver to load the order is modeled as Erlang with mean 12 min and three phases.

Simulate the system to evaluate the improvement that can be attained from the three options relative to the current set-up (so you also need to simulate the current system). Determine a number of replications so that the relative error on each estimate is no more than 5%. Stop each replication at 4 p.m. and do not worry about the trucks remaining in the system at that time, other than recording the value of the number still in the parking lot.

11. Arrivals to the warehouse in Exercise 10 are not stationary. On the book website in `TruckCounts.xls` you will find hourly arrival counts for 30 days. Use these data to estimate arrival rates for a piecewise-constant, nonstationary Poisson arrival process. Implement this arrival process in your warehouse simulation, rerun each case (current system, one forklift, one loading bay, and both a forklift and loading bay) and report the results.
12. Consider the truck arrival data from the previous problem. Let $N(t)$ represent the cumulative number of arrivals by time t . If the process is nonstationary Poisson, then $\text{Var}(N(t))/E(N(t)) = 1$ for all t , or stated differently $\text{Var}(N(t)) = \beta E(N(t))$ with $\beta = 1$. Since you have arrival count data, you can estimate $\text{Var}(N(t))$ and $E(N(t))$ at $t = 1, 2, \dots, 9$ h. Use these data to fit the regression model $\text{Var}(N(t_i)) = \beta E(N(t_i))$ and see if the estimated value of β supports the choice of a nonstationary Poisson arrival process. *Hints:* This is regression through the origin. Also, remember that $N(t_i)$ represents the *total number of arrivals by time t_i* .
13. FirstTreatment, Inc., plans to open for-profit health care clinics throughout the United States Your job is to help FirstTreatment estimate the costs to operate such facilities for different patient loads, which will be part of their business plan. The primary continuing cost for a health care clinic is staffing, which

will be your focus. There are requirements for providing timely treatment, but having more staff and equipment than needed makes a facility expensive to operate. This exercise is based on an example in Kelton et al. (2011).

FirstTreatment's facility design consists of five stations: Sign-In/Triage; Registration; Examination; Trauma; and Treatment. Each of these are staffed separately.

All patients enter at Sign-in/Triage, where a staff member quickly assesses their condition. Patients with severe injuries or illness are sent to Trauma; once their condition is stabilized they move to Treatment; then they are discharged or transported to a hospital.

Patients with less severe problems proceed from Sign-In/Triage to Registration, and then to Examination for evaluation. After the evaluation, roughly 40% are immediately discharged; the others go to Treatment before they are discharged.

The following distributions have been chosen to represent patient contact time:

- **Sign-in/Triage:** Exponential with mean 3 min
- **Registration:** Lognormal with mean 5, and variance 2 min
- **Examination:** Normal with mean 16 and variance 3 min
- **Trauma:** Exponential with mean 90 min
- **Treatment:**
 - For trauma patients, Lognormal with mean 30 and variance 4 min
 - For non-trauma patients, Lognormal with mean 13.3 and variance 2 min

Patient arrivals will be modeled as a Poisson arrival process, and the movement of patients between stations is so brief that it will be treated as 0.

Staffing levels at each station will clearly depend on patient load, percentage of trauma patients, and the patient service-level requirements. FirstTreatment has asked you to look at a range of options for load and trauma (these are distinct scenarios, not random quantities). The percentage of trauma patients depends very much on where the health care facility is located; it could be as low as 8% or as high as 12% of all patients. For load they are guessing as low as 75 to as high 225 patients per day, on average. FirstTreatment clinics will be open from 6 a.m. to 12 p.m. (18 h); once the doors close at midnight patients currently being treated will be completed, but new arrivals will be directed to a local hospital.

Clearly the most important requirement is that trauma patients be stabilized quickly. FirstTreatment has specified that Sign-in/Triage should be "very fast" and that the wait to see a trauma doctor should average less than 5 min. Average delays of 15–20 min are acceptable for other areas of the clinic.

Simulate the health care clinic and provide FirstTreatment with staffing levels required to cover the range of loads for the clinics they might open. Make enough replications so that you can be confident that the service-level requirements have been met. Provide some insight into whether slightly relaxing the service requirements might reduce staffing cost.

14. Arrivals to the clinic in Exercise 13 are not stationary. On the book website in `ClinicCounts.xls` you will find hourly arrival counts for 30 days for a clinic that expects 225 patients/day on average. Use these data to estimate arrival rates for a piecewise-constant, nonstationary Poisson arrival process. Implement this arrival process in your clinic simulation using the staffing level you found for that load. How well does this staffing level do under the more realistic arrival process?
15. Consider the clinic arrival data from the previous exercise. Let $N(t)$ represent the cumulative number of arrivals by time t . If the process is nonstationary Poisson, then $\text{Var}(N(t))/\text{E}(N(t)) = 1$ for all t , or stated differently $\text{Var}(N(t)) = \beta\text{E}(N(t))$ with $\beta = 1$. Since you have arrival count data, you can estimate $\text{Var}(N(t))$ and $\text{E}(N(t))$ at $t = 1, 2, \dots, 18$ h. Use these data to fit the regression model $\text{Var}(N(t_i)) = \beta\text{E}(N(t_i))$ and see if the estimated value of β supports the choice of a nonstationary Poisson arrival process. *Hints:* This is regression through the origin. Also, remember that $N(t_i)$ represents the total number of arrivals by time t_i .
16. Consider the $M/M/\infty$ example of Sect. 7.2.1 used to demonstrate the impact of input uncertainty. Suppose that the arrival rate of the Poisson arrival process was estimated in the following alternative way: A time interval $[0, T]$ was observed and $\tilde{\lambda}$ was estimated by

$$\tilde{\lambda} = \frac{N(T)}{T}$$

where $N(t)$ is the number of arrivals that occurred by time t . This is an unbiased, maximum likelihood estimator of λ if the arrival process is really Poisson. Using this alternative estimator of λ , derive the mean and variance of the simulation estimator described in Sect. 7.2.1.

17. Show that $\text{E}(\hat{\sigma}_I^2) = \sigma_I^2$ under Model (7.11).
18. Consider the stochastic activity network where the time to complete the project is

$$Y = \max\{X_1 + X_4, X_1 + X_3 + X_5, X_2 + X_5\}.$$

A sample of data on activity durations for similar activities can be found on the book website in `SANData.xls`. These data are believed to be exponentially distributed but with different means for each activity, means that can be estimated from the data. The goal is to estimate the expected value of the time to complete the project, $\text{E}(Y)$, using 30 replications. Conduct an experiment to also estimate σ_I^2 , the additional variance due to input uncertainty.

19. Consider the waiting time of the i th customer in a single-server queue as represented by Lindley's equation

$$Y_i = \max\{0, Y_{i-1} + X_{i-1} - A_i\}$$

with $Y_0 = X_0 = 0$. A sample of data on interarrival times and service times can be found on the book website in `MG1Data.xls`. The interarrival times A are believed to be exponentially distributed and the service times X lognormally distributed; the parameters of these distributions can be estimated from

the data. The goal is to estimate the expected value of the time to serve the first 10 customers, $E(\sum_{i=1}^{10} Y_i)$, using 30 replications. Conduct an experiment to also estimate σ_f^2 , the additional variance due to input uncertainty.

Chapter 8

Experiment Design and Analysis

The purpose of simulation, at least in this book, is to estimate the values of performance measures of a stochastic system by conducting a statistical experiment on a computer model of it. This chapter describes principles and methods for the design and analysis of that experiment.

We have considered several different performance measures, including means, probabilities and quantiles, sometimes for systems with a natural time horizon, and sometimes for steady-state systems (as time goes to infinity). For this initial overview, let $\theta(\mathbf{x})$ denote a generic performance measure for scenario \mathbf{x} . A “scenario” defines a specific instance of a more general system. For example, the scenario variable \mathbf{x} could be a vector of cardinal values (e.g., $\mathbf{x}^\top = (15, 9, 6, 3)$ denotes the number Entry Agents and Specialists assigned before and after noon, respectively, in the fax center simulation of Sect. 4.6) or nominal values (e.g., $\mathbf{x} = 2$ denotes the system design for the hospital reception of Sect. 4.3 that employs an electronic kiosk, while $\mathbf{x} = 1$ denotes the current design with a human receptionist). What we call the scenario variable is sometimes called a decision variable, and having it will be useful when we consider comparing system designs to answer questions such as “What is the probability of a special fax being entered in less than 10 minutes when we use a staff of $\mathbf{x}^\top = (15, 9, 6, 3)$ agents?”

The reason for conducting a simulation experiment is to produce an estimator of $\theta(\mathbf{x})$, denoted by $\hat{\theta}(\mathbf{x}; T, n, \mathbf{U})$, which is a function of \mathbf{x} and up to three additional quantities:

Stopping time T : The stopping time is the simulation clock time at which a replication ends; it can be a fixed time (“simulate 8 hours, from 8 AM to 4 PM”) or a random time (“simulate until the stochastic activity network completes” or “simulate until the system fails”). In some contexts T is called runlength. See also the Remark at the end of this section.

Number of replications n : The number of independent and identically distributed replications of the simulation may also be fixed ($n = 30$) or random (“simulate until the width of the confidence interval for $\theta(\mathbf{x})$ is less than 3 minutes”).

(Pseudo)random numbers \mathbf{U} : The estimator $\widehat{\theta}(\mathbf{x}; T, n, \mathbf{U})$ is a random variable because we treat \mathbf{U} as random numbers. Of course, we will actually use pseudorandom numbers; therefore, we may also think of \mathbf{U} as a fixed set of numbers, say \mathbf{u}_1 , that can be reused by fixing the starting seed or stream.

Experiment design consists of selecting the scenarios \mathbf{x} to simulate, controlling the number of replications n to obtain, assigning the pseudorandom numbers \mathbf{U} to drive the simulations and, in steady-state simulation, setting the stopping time T .

Notice that whether or not T is part of the experiment design depends on what we are estimating. For performance measures $\theta(\mathbf{x})$ defined as $T \rightarrow \infty$, a choice as to what finite T to use in the simulation is required. Otherwise, T is typically part of the definition of a replication and is not a design choice. Consider the following examples:

Parking lot ($M(t)/M/\infty$ queue, Sect. 4.2): Here there was only a single scenario so \mathbf{x} is not needed. One performance measure of interest was

$$\theta = \mathbb{E} \left[\max_{0 \leq t \leq 24} N(t) \right]$$

the expected value of the maximum number of cars in the lot during a day. Thus, the stopping time $T = 24$ h is part of the definition of the problem, and the primary design decision is choosing the number of replications n .

Hospital reception ($M/G/1$ queue): Here there were two scenarios, the current human receptionist $\mathbf{x} = 1$ and the electronic kiosk $\mathbf{x} = 2$. The performance measure was

$$\theta(\mathbf{x}) = \mathbb{E}[Y(\mathbf{x})]$$

where $Y(\mathbf{x})$ is the steady-state waiting time before service under scenario \mathbf{x} . In this case both T and n are design variables, and we may even want a design that specifies $n = 1$ and T very large, such as $T = 10^6$ customers.¹ In addition, if we are interested in $\theta(2) - \theta(1)$, the difference in expected waiting times, then we show in Sect. 8.4.1 that there are statistical advantages to assigning both simulations the same set of pseudorandom numbers, say \mathbf{u}_1 . Thus, we could represent the estimator of the difference as

$$\widehat{\theta}(\mathbf{x} = 2; T = 10^6, n = 1, \mathbf{U} = \mathbf{u}_1) - \widehat{\theta}(\mathbf{x} = 1; T = 10^6, n = 1, \mathbf{U} = \mathbf{u}_1).$$

Fax center: Here the scenarios were defined by the number of regular and special agents before and after noon, so $\mathbf{x}^\top = (x_1, x_2, x_3, x_4)$ and there are many possible combinations. If $\mathbf{c}^\top = (c_1, c_2, c_3, c_4)$ is a vector of costs/agent in each of these categories, then the goal of the simulation is

¹ Here we have chosen to measure “time” by the number of customers served, rather than via the simulation clock time. Equivalently, we could let T be the simulation time required to serve 10^6 customers.

$$\begin{aligned} & \min \mathbf{c}^\top \mathbf{x} \\ & \text{subject to:} \\ & \theta_1(\mathbf{x}) \geq 0.96 \\ & \theta_2(\mathbf{x}) \geq 0.80 \end{aligned}$$

where $\theta_1(\mathbf{x})$ and $\theta_2(\mathbf{x})$ are the probabilities that regular and special faxes, respectively, are entered within 10 min of receipt under agent assignment \mathbf{x} . The stopping time T is the time when the last fax arriving before 4 p.m. is completed, and is therefore part of the definition of the problem. The number of replications n might be chosen adaptively depending on the variability of the estimators.

Although it is not possible to entirely separate the topics, Sect. 8.1 covers setting the number of replications, n ; Sect. 8.2 addresses the choice of stopping time for steady-state simulations, T ; and Sect. 8.3 considers which scenarios \mathbf{x} to simulate and the assignment of pseudorandom numbers, \mathbf{U} . Finally, Sect. 8.5 examines a different, but related topic: choice of the estimator, $\hat{\theta}$.

Remark 8.1. Although not critical for understanding design and analysis, we use the term *stopping time* in the formal sense of a stopping time of a stochastic process (see, e.g., Karlin and Taylor 1975). Informally, a stopping time is a point in simulated time, possibly defined by a set of conditions, that can be recognized when it is reached without having to look into the future beyond it. “Simulate for 10 hours of simulation time” defines a stopping time, as does “simulate until 1000 customers have been served” or “simulate until the first fax takes more than 10 minutes to enter.” However, “simulate until you reach the maximum number of cars that will be in the garage” is not a stopping time because we have to simulate the entire 24 h to be sure when the maximum occurs.

8.1 Controlling Error via Replications

Classical statistical analysis was developed to infer characteristics of a population (e.g., income of adults living in England) from a small sample of them, and to infer characteristics of a physical entity or process (e.g., yield per acre of corn hybrids) from an expensive or time-consuming scientific experiment. See, for instance, the exceptional history of statistics by Stigler (1986). In both situations data are expensive or time-consuming to obtain, so it is natural to think of the number of observations as small, fixed, and obtained in a one-time study. This leads to an emphasis on tools that extract as much information as possible from little data, and on methods for measuring the remaining estimation error.

Simulation replications, on the other hand, can often be obtained quickly and inexpensively, and since the experiment is carried out on a computer they can also be obtained sequentially rather than in one trial. This can and should lead to designing experiments to *control error* rather than just measure it. We discuss that approach here.

Our goal in this section is to estimate a performance parameter for a single scenario, so we drop the dependence on \mathbf{x} and denote the parameter by θ . Further, we suppose T is defined by the problem, and since pseudorandom number assignment is only a design decision when we are comparing scenarios, we also drop the dependence on \mathbf{U} and T . Therefore, the goal is to estimate $\theta = \mathbb{E}[\hat{\theta}(n)]$. Further, we consider estimators that are averages across n replications, so to emphasize this we replace the notation $\hat{\theta}(n)$ by $\bar{\theta}(n)$, which is defined to be

$$\bar{\theta}(n) = \frac{1}{n} \sum_{j=1}^n \hat{\theta}_j$$

where $\hat{\theta}_j$ is the unbiased estimator of θ from the j th i.i.d. replication. As an illustration, consider the parking lot model of Sect. 4.2 where

$$\theta = \mathbb{E} \left[\max_{0 \leq t \leq 24} N(t) \right].$$

Then $\hat{\theta}_j$ is the maximum number of cars observed in the parking lot during the j th replication of 24 h.

The normal-theory $(1 - \alpha)100\%$ confidence interval (CI) for θ when n is a fixed number of replications is

$$\bar{\theta}(n) \pm t_{1-\alpha/2, n-1} \frac{S(n)}{\sqrt{n}} \quad (8.1)$$

where

$$S^2(n) = \frac{1}{n-1} \sum_{j=1}^n \left(\hat{\theta}_j - \bar{\theta}(n) \right)^2$$

is the sample variance of the replication estimators. Even when $\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_n$ are not normally distributed, the central limit theorem for averages justifies (8.1) for a large number of replications n because $\sqrt{n}(\bar{\theta}(n) - \theta)$ is asymptotically normal as $n \rightarrow \infty$; see Sect. 5.2.2. When we simulated the parking lot for $n = 1,000$ replications we got a 95% CI of 237.2 ± 0.7 cars.

One way to think about what a CI does is that it provides an error bound on $|\bar{\theta}(n) - \theta|$ that holds with high confidence; specifically

$$\Pr \left\{ |\bar{\theta}(n) - \theta| \leq t_{1-\alpha/2, n-1} \frac{S(n)}{\sqrt{n}} \right\} \approx 1 - \alpha.$$

But better than just measuring error is controlling it, and specifically reducing it below an acceptable threshold for the decision problem at hand. This is different from increasing the confidence with which the bound holds, which is done through our choice of $1 - \alpha$. Instead, we are tightening the bound itself by the choice of n .

The acceptable level of error can be characterized in at least two ways: *Absolute error* means

$$|\bar{\theta}(n) - \theta| \leq \varepsilon \quad (8.2)$$

for given ε , while *relative error* means

$$|\bar{\theta}(n) - \theta| \leq \kappa|\theta| \quad (8.3)$$

for given $0 < \kappa < 1$. The absolute error ε is in the same units as the performance measure θ (minutes, dollars, cars). For the parking lot simulation we might consider an absolute error of $\varepsilon = 1$ car. On the other hand, κ is a fraction of θ itself. Relative error is often stated as a percentage; for instance $\kappa = 0.01$ is “1% relative error.” For the parking lot simulation it appears that 1% relative error is about 2.4 cars.

An intuitively appealing approach to achieve a given absolute error is to make N replications, where

$$N = \min \{n : t_{1-\alpha/2, n-1} S(n) / \sqrt{n} \leq \varepsilon\}. \quad (8.4)$$

Similarly, to achieve given relative error make N replications where

$$N = \min \{n : t_{1-\alpha/2, n-1} S(n) / \sqrt{n} \leq \kappa|\bar{\theta}(n)|\}. \quad (8.5)$$

Notice that in both cases N is a random variable, so the validity of this approach is not supported by the usual central limit theorem for averages. Instead, theoretical support comes from central limit theorems for $\sqrt{N}(\bar{\theta}(N) - \theta)$ as $\varepsilon \rightarrow 0$ and $\kappa \rightarrow 0$, respectively. The classical references are Chow and Robbins (1965) and Nádas (1969). A paper that specifically addresses the simulation setting is Glynn and Whitt (1992); see also Alexopoulos (2006). Intuitively, this works because driving the absolute error ε or the relative error κ to zero drives N to infinity, but the actual proof is more subtle.

However, it is important to understand that, unlike the t -distribution CI (8.1), even when $\hat{\theta}_1, \hat{\theta}_2, \dots$ are normally distributed the bounds (8.2) and (8.3) are not guaranteed to hold with probability $1 - \alpha$ when N is chosen in this way. Instead, as we become more demanding (accept smaller and smaller error) these stopping rules are valid.

To implement (8.4) or (8.5) we need to be able to update $\bar{\theta}(n)$ and $S^2(n)$ efficiently as we increase the number of replications. Two relationships help:

$$\bar{\theta}(n) = \bar{\theta}(n-1) + \frac{\hat{\theta}_n - \bar{\theta}(n-1)}{n} \quad (8.6)$$

$$\sum_{j=1}^n (\hat{\theta}_j - \bar{\theta}(n))^2 = \sum_{j=1}^{n-1} (\hat{\theta}_j - \bar{\theta}(n-1))^2 + (\hat{\theta}_n - \bar{\theta}(n)) (\hat{\theta}_n - \bar{\theta}(n-1)). \quad (8.7)$$

The following algorithm, due to Knuth (1998), implements this recursion.

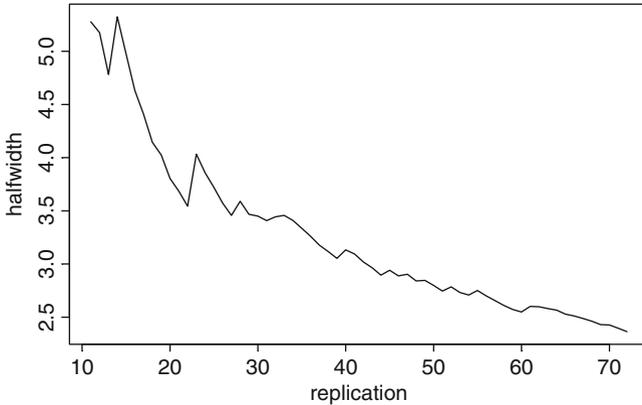


Fig. 8.1 Confidence interval halfwidth as a function of the number of replications for the parking lot simulation with 1% relative error stopping rule

1. $n = 0$, $\text{mean} = 0$ and $\text{sum} = 0$
2. Obtain replication result $\hat{\theta}_{n+1}$
 - a. $n = n + 1$
 - b. $\text{diff} = \hat{\theta}_n - \text{mean}$
 - c. $\text{mean} = \text{mean} + \text{diff}/n$
 - d. $\text{sum} = \text{sum} + \text{diff} * (\hat{\theta}_n - \text{mean})$
 - e. If $n > 1$ then $\bar{\theta}(n) = \text{mean}$ and $S^2(n) = \text{sum}/(n - 1)$
3. Next replication

The particular stopping test (8.4) or (8.5) would be inserted in Step 2e, and N is the first n for which it is satisfied. However, there are good reasons not to implement the algorithm exactly that way.

Unfortunately, rules such as (8.4) or (8.5) are subject to stopping too early, before N has a chance to become large. This is because the sample variance $S^2(n)$ can be quite variable for small n , meaning there is a chance that it will be far smaller than the true variance; this leads to premature termination and a CI that is too short. Unless simulation replications are very time-consuming, the following simple guideline will help: Make at least 10 replications before applying the stopping test for the first time, with 60 replications being better. A further advantage of starting at $n = 60$ is that the t quantiles $t_{1-\alpha/2, n-1} \approx z_{1-\alpha/2}$ for $n \geq 60$, so there is no need to obtain the t values.

Figure 8.1 is a plot of the CI halfwidth against the number of replications for the parking lot problem with the 1% relative error stopping rule, computing the confidence interval for the first time at $n = 11$ replications. The simulation stops after $N = 72$ replications with confidence interval 236.306 ± 2.363 cars. Thus, the estimated relative error is $2.363/236.306 \approx 1\%$. Notice that due to sampling variability

the halfwidths do not decrease monotonically, and it is this variability that can cause early stopping.

One type of performance measure that does not fit neatly into this framework is a quantile. As discussed in Sect. 7.1.1 the q -quantile of a continuous-valued random variable Y is $\vartheta = F_Y^{-1}(q)$, and the standard estimator is $\hat{\vartheta} = Y_{(\lceil nq \rceil)}$, where $Y_{(1)} \leq Y_{(2)} \leq \dots \leq Y_{(n)}$ are the sorted values from all n replications. Therefore, we need the results from all n replications to form a single quantile estimator. The same is true of using a sample variance to estimate the $\text{Var}(Y)$.

Because we can often afford to make the number of replications very large, we can adapt the methods in this section to performance measures such as quantiles and variances using an approach called *batching*. The idea is that we redefine the j th output to be a statistic computed from a batch of b individual replications; this statistic can be an order statistic, a sample variance, or any other statistic of interest:

$$\underbrace{Y_1, \dots, Y_b}_{\hat{\theta}_1}, \underbrace{Y_{b+1}, \dots, Y_{2b}}_{\hat{\theta}_2}, \dots, \underbrace{Y_{(n-1)b+1}, \dots, Y_{nb}}_{\hat{\theta}_n}.$$

With this adjustment the methods described in this section can be applied directly to the batch statistics. However, a batch size b needs to be chosen, and in the case of quantile estimation this is particularly critical because quantile estimators are biased, with that bias decreasing in b . Thus, b should not be small.

8.2 Design and Analysis for Steady-State Simulation

In this section we address the challenging problem of steady-state simulation. To do so we will consider the simulation of a single scenario, so we can drop the scenario variable \mathbf{x} and the random-number assignment \mathbf{U} from the notation. Therefore, the goal is to estimate a steady-state parameter θ with an estimator $\hat{\theta}(T, n)$, since the number of replications n and the run length T are both design choices. We will also introduce an additional design decision that is only relevant for steady-state simulation, an amount of data to delete, d .

For most of the section we will focus on the problem of estimating the steady-state mean μ from a discrete-time output series Y_1, Y_2, \dots, Y_m , so the run length is the number of observations $T = m$ and the estimator will be a sample mean. Therefore, in this section we employ the notation

$$\hat{\theta}(T, n) = \bar{Y}(n, m, d)$$

as defined below.

We used the $M/G/1$ queue (see Sect. 3.2) as an example of a steady-state simulation, one for which the quantity we want to estimate is defined in the limit as the simulation run length goes to infinity. For the $M/G/1$ queue this suggested that

we make a very long simulation run (large number of customers m) and estimate, say, the steady-state mean waiting time μ by the average of the observed waiting times Y_1, Y_2, \dots, Y_m . But we noticed that the waiting times early in the run tend to be smaller than μ because the queue starts empty, which causes $\bar{Y}(m)$ to be biased low. To reduce this effect, we suggested letting the simulation generate waiting times for a while (say d of them) before starting to actually include them in the sample average. That is, we used the truncated sample average

$$\bar{Y}(m, d) = \frac{1}{m-d} \sum_{i=d+1}^m Y_i. \quad (8.8)$$

We may also choose to make n replications, yielding n i.i.d. averages $\bar{Y}_1(m, d), \bar{Y}_2(m, d), \dots, \bar{Y}_n(m, d)$, giving the overall average

$$\bar{Y}(n, m, d) = \frac{1}{n} \sum_{j=1}^n \bar{Y}_j(m, d). \quad (8.9)$$

Of course, this includes the special case of $n = 1$ replication. Experiment design and analysis for steady-state simulation focuses on choosing n, m and d . We now give two precise statements of the “steady-state simulation problem.” Neither problem is directly solvable in practice, but knowing which one best represents your objective helps determine the approach you will use.

Fixed-precision problem: Given $\varepsilon > 0$ or $0 < \kappa < 1$, choose n, m and d so that $\sqrt{\text{MSE}(\bar{Y}(n, m, d))} \leq \varepsilon$ (absolute error), or $\sqrt{\text{MSE}(\bar{Y}(n, m, d))}/\mu \leq \kappa$ (relative error). That is, choose n, m and d so that a fixed precision is achieved. This statement of the problem favors large d to effectively eliminate bias, and $n > 1$ replications to make it easier to quantify error via, say, a confidence interval.

Fixed-budget problem: Given a budget of N observations, choose n, m and d to minimize $\text{MSE}(\bar{Y}(n, m, d))$ subject to $nm \leq N$ and $d < m$. That is, spend a fixed amount of simulation effort in a way that minimizes MSE. Examination of the asymptotic MSE (8.11) or the AR(1) result (8.12) below makes it immediately obvious that this formulation favors solutions with $n = 1$ replication and $m = N$.

Notice that both formulations focus on $\text{MSE}(\bar{Y}(n, m, d))$ as the measure of the quality of the estimator, which makes sense because bias is a feature of steady-state simulation. The next two sections present methods for attacking each of these problems, but to have them make sense we need to understand the properties of the Estimator (8.9).

We argued in Sect. 5.2.3 that when steady-state makes sense then we should expect that the mean squared error of the sample average of one long replication can be approximated by

$$\text{MSE}(\bar{Y}(1, m, 0)) \approx \frac{\beta^2}{m^2} + \frac{\gamma^2}{m} \quad (8.10)$$

where β is the asymptotic bias and γ^2 is the asymptotic variance. We referred to this quantity as the asymptotic MSE. We can also argue that for large m and d much smaller than m that

$$\text{MSE}(\bar{Y}(1, m, d)) \approx \frac{\beta(d)^2}{(m-d)^2} + \frac{\gamma^2}{m-d}$$

where $\beta(0) = \beta$. Why should the asymptotic bias be a function of the amount of data deleted, d ? Recall that

$$\beta = \sum_{i=1}^{\infty} (\mathbb{E}(Y_i) - \mu)$$

which can be thought of as the area under the bias curve from 0 to ∞ , or the total accumulated bias. And since the bias goes to 0 as $1/m$, we approximated the bias of $\bar{Y}(1, m, 0)$ as β/m . Similarly, let

$$\beta(d) = \sum_{i=d+1}^{\infty} (\mathbb{E}(Y_i) - \mu)$$

be the total accumulated bias from $d+1$ to ∞ , so the bias of $\bar{Y}(1, m, d)$ is approximately $\beta(d)/(m-d)$. Adjusting the asymptotic bias for deletion is important because we expect the bias to be largest early in the replication.

Finally, for general number of replications n ,

$$\text{MSE}(\bar{Y}(n, m, d)) \approx \frac{\beta(d)^2}{(m-d)^2} + \frac{\gamma^2}{n(m-d)}. \quad (8.11)$$

That the variance decreases as $1/n$ is expected. That the bias is unaffected by replications is a consequence of

$$\mathbb{E}(\bar{Y}(n, m, d) - \mu) = \frac{1}{n} \sum_{j=1}^n \mathbb{E}(\bar{Y}_j(m, d) - \mu) = \mathbb{E}(\bar{Y}_1(m, d) - \mu).$$

Replications reduce variance, but not bias, in steady-state simulation.

As a specific example, Exercise 2 asks you to show that the asymptotic MSE of the AR(1) surrogate process is

$$\text{MSE}(\bar{Y}(n, m, d)) \approx \frac{(y_0 - \mu)^2 \varphi^{2d+2}}{(m-d)^2 (1-\varphi)^2} + \frac{\sigma^2}{n(m-d)(1-\varphi)^2} \quad (8.12)$$

and further to prove that for $\varphi > 0$ the asymptotic squared bias (the first term) is decreasing in d for m large enough. Clearly, the variance $\sigma^2 / (n(m-d)(1-\varphi)^2)$ is increasing in d . Thus, there is a bias-variance trade off in choosing a deletion point d for steady-state simulation.

Expression (8.12) illustrates features of the steady-state simulation problem that affect the MSE: Starting far from steady-state conditions (as represented by

$(y_0 - \mu)^2$ being large), or having slowly diminishing bias (as represented by $|\phi|$ close to one) make the bias the dominant feature of MSE, suggesting deletion. A large variance of the outputs (as represented by σ^2 large) makes the variance dominant, suggesting a long run length m or number of replications n .

8.2.1 Fixed-Precision Problem

The fixed-precision setting makes sense when the simulation budget is not a constraint: typically this means that the simulation model executes fast enough, and that the time until the results are needed is long enough, that we can afford to “simulate until we are done.” This is often the case in practice, and it makes no sense to enforce an arbitrary simulation budget when one does not exist.

The approach we take to solve this problem consists of two steps:

1. Determine a run length m and deletion point \hat{d} such that the bias $\beta(\hat{d})/(m - \hat{d})$ is effectively 0. Employ this deletion point to obtain truncated averages that are (effectively) without bias.
2. Measure the remaining statistical error via a confidence interval or standard error, and increase the run length or number of replications until the desired precision has been achieved.

The specific methods applied to achieve Steps 1 and 2 depend on whether we plan to make $n = 1$ or $n > 1$ replications. We focus on the $n > 1$ case; methods described in Sect. 8.2.2 for a fixed budget could be adapted to obtain fixed precision from a single run by letting batches replace replications.

The primary feature that makes steady-state simulation hard is that unlike variance, the bias $\beta(d)/(m - d)$ is not directly estimable. However, a sufficient condition for $\beta(d)/(m - d) \approx 0$ is that $E(Y_i) - \mu \approx 0$ for all $i > d$, and this will be true if $E(Y_i)$ is essentially unchanging for $i > d$. Even though the bias cannot be estimated, the change in $E(Y_i)$ as a function of i can, particularly if we can afford to make replications. Therefore, we solve Step 1 by estimating the $E(Y_i)$ vs. i curve and observing where it appears to stop changing.

8.2.1.1 Estimating the Deletion Point

1. Obtain output data $Y_{ij}, i = 1, 2, \dots, m$ and $j = 1, 2, \dots, n$. For instance, in the $M/G/1$ simulation $Y_{3,8}$ is the waiting time of the third customer on the eighth replication. The run length m will be a guess initially, while the number of replications n should be at least 10.
2. Average the i th observation across all n replications:

$$\bar{Y}_i = \frac{1}{n} \sum_{j=1}^n Y_{ij}$$

for $i = 1, 2, \dots, m$. For instance, in the $M/G/1$ simulation \bar{Y}_3 would be the average waiting time of the third customer to arrive in each of the n replications.

3. Examine or plot \bar{Y}_i vs. i , or perhaps a smoothed version of it.
 - (a) If the data are too variable to detect the trend, then increase the number of replications n .
 - (b) If the plot still seems to be trending up or down, then increase the run length m .
 - (c) If there is an observation number \hat{d} such that beyond it the plot seems to be varying consistently around a fixed central value, then choose \hat{d} as the deletion point.

If done successfully, then we have selected \hat{d} such that

$$\text{MSE}(\bar{Y}_j(m, \hat{d})) \approx 0 + \frac{\gamma^2}{m - \hat{d}}$$

and therefore

$$\text{MSE}(\bar{Y}(n, m, \hat{d})) \approx \text{Var}(\bar{Y}(n, m, \hat{d})) \approx \frac{\gamma^2}{n(m - \hat{d})}.$$

To solve Step 2, we form a confidence interval for μ

$$\bar{Y}(n, m, \hat{d}) \pm t_{1-\alpha/2, n-1} \frac{S(n, m, \hat{d})}{\sqrt{n}}$$

where

$$S^2(n, m, \hat{d}) = \frac{1}{n-1} \sum_{j=1}^n (\bar{Y}_j(m, \hat{d}) - \bar{Y}(n, m, \hat{d}))^2$$

is the sample variance of the replication averages. The degrees of freedom analysis presented in the next section suggests that if the confidence interval is not yet short enough to meet the precision requirement, then it makes sense to increase the number of replications if $n < 30$, and otherwise to increase the run length m . However, it is usually more convenient to fix m and just increase the number of replications until the confidence interval is short enough. If this approach is adopted, then make sure that m is substantially larger than d , say $m \approx 10d$.

To achieve fixed absolute error we increase the number of replications n until

$$t_{1-\alpha/2, n-1} S(n, m, \hat{d}) / \sqrt{n} \leq \varepsilon.$$

Similarly, to achieve fixed relative error we increase the number of replications until

$$\frac{t_{1-\alpha/2, n-1} S(n, m, \hat{d}) / \sqrt{n}}{\bar{Y}(n, m, \hat{d})} \leq \kappa.$$

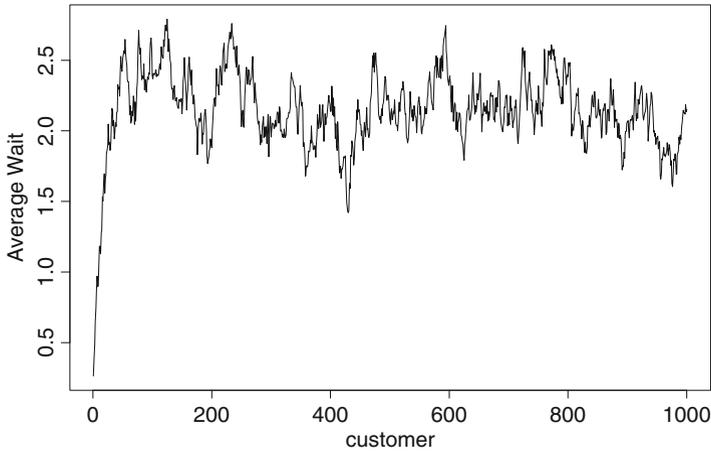


Fig. 8.2 Mean plot, by customer number, averaged across $n = 100$ replications of the $M/G/1$ queue

To illustrate these ideas, recall the $M/G/1$ queueing simulation of Sect. 4.3, where the goal was to estimate the steady-state mean customer waiting time in queue. Suppose we desired this estimate to have no more than 3% relative error.

Figure 8.2 shows the average waiting time, by customer number, for the first $m = 1,000$ customers averaged across $n = 100$ replications. The bias due to starting the simulation empty is apparent. At somewhere around 100 customers the average waiting times seem to vary around a central value, so we might take $\hat{d} = 100$. This is a subjective judgement; additional replications, or smoothing of the plot, would be helpful to confirm this choice, and there is no penalty (other than computation time) for making it larger.

Given $\hat{d} = 100$, a convenient choice for the run length m is 1,100 customers (at least ten times as large). Thus, each replication will generate waiting times for the first $m = 1,100$ customers, and report the average of last $m - \hat{d} = 1,000$ of them. A guess of $n = 500$ replications satisfied the relative error requirement, as it gave

$$\frac{t_{0.975,499}S(500,1100,100)/\sqrt{500}}{\bar{Y}(500,1100,100)} \approx \frac{0.053}{2.154} \approx 0.025 < 0.03.$$

Remark 8.2. This $M/G/1$ queue, which represents hospital reception, is not a very congested system, so the amount of data deletion required is very small. This is fine for a book example, but should not be considered representative of all simulations, or even all queueing simulations. Heavily congested queueing systems may require substantially more data deletion.

8.2.2 Fixed-Budget Problem

When we have a fixed budget of N observations, and no sense whether it is generous or tight, the natural experiment design is to make a single replication ($n = 1, m = N$) to have the best chance of overcoming the bias. We will again have two steps: selection of a deletion point d , and estimation of the remaining error in the point estimator.

8.2.2.1 Deletion with a Fixed Budget

We consider three options for choosing a deletion point d :

No deletion: As shown in Eq. (8.10), the bias component of MSE goes down as $1/m^2$, while variance only diminishes as $1/m$. With no direct information about the bias, deleting no data is a viable option.

Use system knowledge: Although we can treat steady-state simulation as an abstract statistical problem, in applications we are simulating a real or conceptual system that we understand well enough to model it. Insight about how long it might take the real system to “warm up” can be used to set a deletion point. If a factory starts empty, how many hours or days of production would be required to ramp up to normal operations? If all echelons of a supply chain are fully stocked, how many weeks would it take for orders, transportation and inventory levels to start behaving normally? Answers to these questions provide deletion points.

Data-driven deletion point: A number of methods have been proposed for choosing a deletion point based on the output data. These include the following. Note that the first two approaches try to eliminate bias, while the third tries to actually minimize the MSE of the resulting estimator.

- Plot the cumulative sample mean: Although we cannot average across replications to reveal the bias trend, we can plot $\sum_{i=1}^t Y_i/t$ vs. t , for $t = 1, 2, \dots, m$, and look for a point at which it stabilizes around a central value. This is often done because it is easy, but it tends to delete many more observations than necessary since the cumulative average retains all of the most biased observations. In a fixed-budget setting it is not typically a good idea to be overly conservative with respect to bias since it results in a corresponding penalty in terms of estimator variance. Schruben (1982) provides a single-replication plot that is more sensitive to bias, but is less directly useful for determining a deletion point.
- Test for bias: If a deletion point d can be chosen based on knowledge of the system, or even a guess, then statistical tests have been proposed in which the null hypothesis is $H_0 : E(Y_{d+1}) = \dots = E(Y_m)$. See, for instance, Schruben et al. (1983). A description of a number of these tests, and evaluation of their performance, can be found in Cash et al. (1992).
- Estimate the MSE-optimal deletion point: Although the MSE as a function of d cannot be estimated directly, the marginal standard error rule (MSER, White

1997) minimizes the value of a statistic whose expectation is asymptotically proportional to MSE (Pasupathy and Schmeiser 2010). Since it attacks the stated problem directly, and is easy to implement, we describe MSER in more detail below.

The $\text{MSER}(d)$ statistic is

$$\text{MSER}(d) = \frac{1}{(m-d)^2} \sum_{i=d+1}^m (Y_i - \bar{Y}(m,d))^2 \quad (8.13)$$

and the estimated deletion point is chosen to minimize $\text{MSER}(d)$. Specific rules include

$$\hat{d} = \operatorname{argmin}_{d=0, \dots, \lfloor m/2 \rfloor} \text{MSER}(d) \quad (8.14)$$

and

$$\tilde{d} = \min \{d : \text{MSER}(d) \leq \min[\text{MSER}(d-1), \text{MSER}(d+1)]\}. \quad (8.15)$$

The choice \hat{d} minimizes MSER restricted to the first half of the output series, while \tilde{d} yields the first local minimum of MSER; both choices recognize that the MSER statistic is quite variable when d is close to m and try to avoid choosing d too large just because of output variability.

Clearly the bias and variance of the outputs Y_1, Y_2, \dots, Y_m influence the value of MSER, but it is not immediately obvious how minimizing MSER relates to minimizing MSE. Intuitively, either a strong trend in the mean of the series Y_1, Y_2, \dots, Y_m (which suggests that the amount of deletion should be large) or substantial marginal variance $\text{Var}(Y_i)$ (which suggests that the amount of deletion should be small) will cause (8.13) to be large, so minimizing MSER balances these contributions. Pasupathy and Schmeiser (2010) showed that, under very general conditions, for every $d = 0, 1, 2, \dots$

$$\lim_{m \rightarrow \infty} \frac{\text{MSE}(\bar{Y}(m,d))}{\text{E}(\text{MSER}(d))} = c$$

where c is a constant that depends on the output process. That is, for large m , MSE and MSER are proportional, in expectation, so that minimizing the expected value of MSER is equivalent to minimizing MSE. Of course, the expected value of MSER is not known, so we minimize the statistic instead.

Exercise 5 asks you to show that

$$\sum_{i=d+1}^m (Y_i - \bar{Y}(m,d))^2 = \sum_{i=d+1}^m Y_i^2 - \frac{1}{m-d} \left(\sum_{i=d+1}^m Y_i \right)^2.$$

Using this formula, the MSER statistic can be computed in one pass through the data by starting from the end of the series and working backward to the beginning:

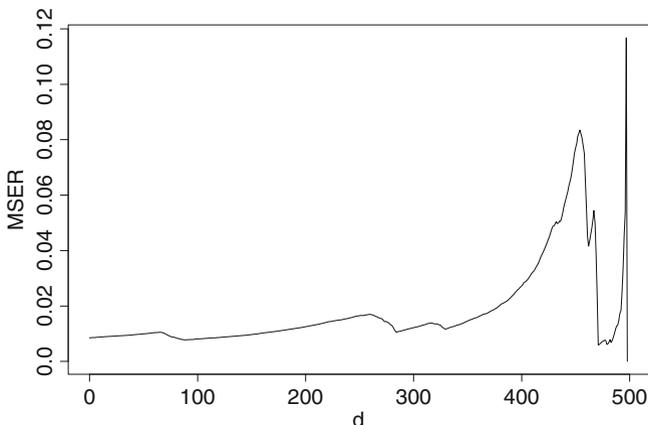


Fig. 8.3 The $MSER(d)$ statistic for the $M/G/1$ example with run length $m = 500$ customers. In this example $\hat{d} = \tilde{d} = 88$

1. Set $s = 0, q = 0$
2. For $d = m - 1$ to 0
 - a. $s = s + Y_{d+1}$
 - b. $q = q + Y_{d+1}^2$
 - c. $MSER(d) = (q - s^2 / (m - d)) / (m - d)^2$
3. Next d

Figure 8.3 shows a plot of $MSER(d)$ for one replication of $m = 500$ waiting times from the $M/G/1$ simulation of Sect. 4.3. In this case $\hat{d} = \tilde{d} = 88$, which suggests that estimating the steady-state waiting time by averaging the last $500 - 88 = 412$ values will minimize the MSE of the estimate.

8.2.2.2 Error Estimation with a Fixed Budget

When we had n replications, we exploited the fact that

$$\text{Var}(\bar{Y}(n, m, d)) = \frac{\text{Var}(\bar{Y}(1, m, d))}{n}$$

because replications allowed us to estimate $\text{Var}(\bar{Y}(1, m, d))$. The relationship between the variance of the sample mean and the variance of the component replications is one of the most important in all of statistics.

When we only have a single replication, Y_1, Y_2, \dots, Y_m (perhaps after some deletion), we can exploit a similar relationship, provided m is sufficiently large: Let $\bar{Y}(m)$ be the overall sample mean (dropping d and n to simplify notation). Then we

know that $\text{Var}(\bar{Y}(m)) \approx \gamma^2/m$. Here is the key insight: For $b < m$, but still large, $\text{Var}(\bar{Y}(b)) \approx \gamma^2/b$ as well. Therefore, if both m and b are large,

$$\text{Var}(\bar{Y}(m)) \approx \frac{b}{m} \text{Var}(\bar{Y}(b)). \quad (8.16)$$

However, rather than making replications, we can estimate $\text{Var}(\bar{Y}(b))$ by forming $k = m/b$ batch means as shown below:

$$\underbrace{Y, \dots, Y}_{\text{deleted}} \underbrace{Y_1, \dots, Y_b}_{\bar{Y}_1(b)} \underbrace{Y_{b+1}, \dots, Y_{2b}}_{\bar{Y}_2(b)} \dots \underbrace{Y_{(k-1)b+1}, \dots, Y_{kb}}_{\bar{Y}_k(b)}.$$

We refer to b as the batch size, and k as the number of batches. The natural estimator of $\text{Var}(\bar{Y}(b))$ is then the sample variance of the batch means

$$S^2(k) = \frac{1}{k-1} \sum_{j=1}^k (\bar{Y}_j(b) - \bar{Y}(m))^2.$$

This yields the batch means variance estimator $\widehat{\text{Var}}(\bar{Y}(m)) = (b/m)S^2(k)$, or equivalently the batch means estimator of the asymptotic variance constant $\widehat{\gamma}^2 = bS^2(k)$. Finally, an approximate confidence interval for μ is

$$\bar{Y}(m) \pm t_{1-\alpha/2, k-1} \frac{\widehat{\gamma}}{\sqrt{m}}$$

which is algebraically equivalent to

$$\bar{Y}(m) \pm t_{1-\alpha/2, k-1} \frac{S(k)}{\sqrt{k}}. \quad (8.17)$$

The development above made use of a lot of “ \approx ” arguments. When might this be a good approximation? Recall from Eq. (5.5) that

$$\gamma^2 = \sigma^2 \left(1 + 2 \sum_{i=1}^{\infty} \rho_i \right).$$

Because the correlations have to be summable for γ^2 to exist, we expect the autocorrelations to diminish with i , and to be effectively 0 for $i > b^*$, for some b^* . Therefore, $\sigma^2 \left(1 + 2 \sum_{i=1}^{b^*} \rho_i \right)$ is a really good approximation of γ^2 . If the batch size $b \geq b^*$ (which of course means $m > b^*$ as well), then the relationship (8.16) holds. Thus, the batch size needs to capture the correlation structure of the output process.

For the confidence interval (8.17) to be asymptotically valid requires a bit more: As the sample size m goes to infinity with the number of batches k fixed (so that the batch size is also increasing), the batch means $\bar{Y}_1(b), \bar{Y}_2(b), \dots, \bar{Y}_k(b)$ need to

become independent and normally distributed, which can be shown to be true under relatively mild conditions (e.g., Steiger and Wilson 2001).

The problem of deciding when b is large enough is a difficult one. Obviously the larger b is the more likely (8.16) is to hold. A number of data-driven algorithms for selecting batch size have been published, with the goal of either delivering a valid confidence interval or standard error estimate at the end. Some algorithms also control the run length to achieve a fixed precision, while others assume the run length is given. An example, with references to the broader literature, is Tafazzoli and Wilson (2011).

The data driven batching algorithms are sometimes complicated, meaning that implementation could involve substantial effort. If the simulation budget N is really fixed and cannot be extended, then the following procedure is sensible and relatively simple to apply:

1. Make a single replication of length $m = N$.
2. Apply MSER to obtain a deletion point d (or for an even simpler procedure, set $d = 0$).
3. Divide the remaining $m - d$ observations into from $10 \leq k \leq 30$ batches, looking for a value of k that divides $m - d$ close to evenly (if there are data left over, delete from the beginning).
4. Compute the sample mean and form the batch means confidence interval (8.17).

The assumption behind this approach is that $m = N$ is very large, large enough that it is possible to estimate μ precisely and form an approximately valid confidence interval (and possibly even large enough that data deletion is not necessary). When this is the case, the analysis in Schmeiser (1982) shows that little is lost by using a relatively small number of batches $10 \leq k \leq 30$, while on the other hand there is risk if the number of batches k is too large making the batch size $b < b^*$ (too small).

The reason that little is lost provided $k \geq 10$ can be understood by observing the impact of number of batches on (8.17). Notice that the confidence interval depends on the estimator $\hat{\gamma}^2$ of the asymptotic variance constant, γ^2 , and the batch means variance estimator typically will be valid for any value of k small enough that $b > b^*$.

So why not pick $k = 2$? Consider the ratio of $t_{1-\alpha/2, k-1}$ to $z_{1-\alpha/2} = t_{1-\alpha/2, \infty}$ for $1 - \alpha = 0.95$, shown in Table 8.1. At $k = 2$ it is about 6.5 times larger than it would be if k was very large. However, at $k = 10$ it is only 1.15 times larger, dropping to 1.04 times larger at $k = 30$ batches. Since the risk of an invalid confidence interval increases as k increases, $10 \leq k \leq 30$ is a good working range.

The method described here for estimating γ^2 is usually called the *nonoverlapping batch means method*, and it is one of many that have been proposed. An overview can be found in Goldsman and Nelson (2006).

For the $M/G/1$ simulation of Sect. 4.3 we made a single replication of $m = 30,000$ customer waiting times. The minimum $\text{MSER}(d)$ occurred at $\hat{d} = \tilde{d} = 0$, which is not surprising with a long run for a system that exhibits little bias: apparently the benefit from reducing the (small) bias with 30,000 observations was not substantial relative to the increased variance from deleting data. These data were

Table 8.1 Relative impact of number of batches k on the 0.975 quantile of the t distribution

k	$\frac{t_{0.975,k-1}}{z_{0.975}}$
2	6.48
5	1.42
10	1.15
20	1.07
30	1.04
60	1.02
∞	1.00

then batched into $k = 30$ batches of size $b = 1,000$ waiting times and the sample means computed from each batch. Applying the CI (8.17) yielded the 95% CI 2.24 ± 0.22 min.

What if an error of ± 0.22 min is considered too large? Ideally, the run length m should be increased and the data are rebatched into $10 \leq k \leq 30$ batches. If the raw data were not retained, then additional batches of the same size b should be generated, increasing k .

8.2.3 Batch Statistics

The preceding sections covered estimating the mean of a steady-state simulation output process from discrete-time output data. However, performance measures such as queue length and inventory level are continuous-time outputs, $Y(t)$, $0 \leq t \leq T$. In the fixed-budget, single-replication setting, continuous-time outputs can be converted into discrete-time outputs using the batching transformation

$$Y_i = \frac{1}{\Delta t} \int_{(i-1)\Delta t}^{i\Delta t} Y(t) dt \quad (8.18)$$

for $i = 1, 2, \dots, m$ where $m = \lfloor T/\Delta t \rfloor$ and $\Delta t > 0$ is a time interval. For instance, if $Y(t)$ is queue length, time $t = 0$ corresponds to 8 a.m., and $\Delta t = 15$ min, then Y_7 is the average queue length from 9:30 a.m. (time 90 min) to 9:45 a.m. (time 105 min). Tools such as MSER and batch means can then be applied to Y_1, Y_2, \dots, Y_m . It is important that Δt be large enough that the value of $Y(t)$ is likely to change one or more times during the interval; otherwise the transformed process Y_1, Y_2, \dots, Y_m may exhibit extremely strong positive correlation with the same value appearing two or more times in succession.

As shown at the end of Sect. 8.1 for across-replication data, the concept of batching can be applied to performance measures other than means and this is the case even for a single, long replication of a steady-state simulation provided the performance measure θ can be consistently estimated by a statistic $\hat{\theta}$ that is a function of stationary, but dependent data. The basic outputs from the replication are divided

into k batches of b individual observations, and from the j th batch we compute a $\widehat{\theta}_j(b)$ statistic:

$$\underbrace{Y, \dots, Y}_{\text{deleted}}, \underbrace{Y_1, \dots, Y_b}_{\widehat{\theta}_1(b)}, \underbrace{Y_{b+1}, \dots, Y_{2b}}_{\widehat{\theta}_2(b)}, \dots, \underbrace{Y_{(k-1)b+1}, \dots, Y_{kb}}_{\widehat{\theta}_k(b)}.$$

The performance measure θ is estimated by the average

$$\bar{\theta}(k) = \frac{1}{k} \sum_{j=1}^k \widehat{\theta}_j(b)$$

and not by applying $\widehat{\theta}$ to the entire data set; in the case of batch means the average of the batch means is equal to the overall average, but this is not true for other statistics.

For example, suppose that we want to estimate the marginal variance $\sigma^2 = \text{Var}(Y)$. If Y_1, Y_2, \dots, Y_m is a stationary but possibly dependent output process (say, after deletion) with $\text{Var}(Y_i) = \sigma^2$, then it can be shown under mild conditions that

$$\frac{1}{m-1} \sum_{i=1}^m (Y_i - \bar{Y}(m))^2 \xrightarrow{a.s.} \sigma^2$$

as $m \rightarrow \infty$. That is, it is a consistent estimator even though the output data are dependent. To apply the batching method, the batch statistics are

$$\widehat{\theta}_j(b) = \frac{1}{b-1} \sum_{i=(j-1)b+1}^{jb} (Y_i - \bar{Y}_j(b))^2, \quad j = 1, 2, \dots, k$$

the sample variance from within each batch of observations. The marginal variance σ^2 is estimated by

$$\bar{\theta}(k) = \frac{1}{k} \sum_{j=1}^k \widehat{\theta}_j(b)$$

the average of the batch statistics, and a confidence interval is formed as in (8.17) using the variance estimator

$$S^2(k) = \frac{1}{k-1} \sum_{j=1}^k \left(\widehat{\theta}_j(b) - \bar{\theta}(k) \right)^2.$$

8.2.4 Steady-State Simulation: A Postscript

Two things about the steady-state simulation problem can be stated with confidence: it is a difficult problem to solve, and there have been many, many attempts to solve it. The tools presented in the previous sections were selected because they are easy

to use and they illustrate what we believe are the important concepts. There are, however, other methods that may have significant advantages and new methods are being invented regularly. The best way to research solutions past and present is via the archive of the *Proceedings of the Winter Simulation Conference* at <http://www.wintersim.org/>.

In any event, the important concepts are these:

1. There is a tradeoff between bias and variance, and both matter, as reflected in the MSE.
2. Bias diminishes faster than variance with increasing run length, but does not diminish at all with increasing numbers of replications. Thus, no solution will have a very large number of short replications.
3. How you attack the steady-state simulation problem depends on whether you can afford (usually in terms of time available) to treat it as a fixed-precision or a fixed-budget problem. When data are cheap (i.e., fast) a fixed-precision approach may waste data but it gives more assurance that you get what you want.

8.3 Simulation Optimization

We now focus on the problem of finding a good, or perhaps even the best, scenario \mathbf{x} from a collection of feasible scenarios. We will primarily consider the formulation

$$\begin{aligned} \min \theta(\mathbf{x}) \\ \mathbf{x} \in C \end{aligned} \tag{8.19}$$

where $\theta(\mathbf{x})$ is the performance measure of interest, and C is the feasible set or region for the d -dimensional scenario variable \mathbf{x} . The feasible region C is often defined by a collection of constraints on \mathbf{x} , but it may also be just a list of options. What makes this a *simulation optimization* (SO) problem is the need to estimate the scenario performance measure $\theta(\mathbf{x})$ using a simulation-based estimator $\hat{\theta}(\mathbf{x}; T, n, \mathbf{U})$. In Sect. 8.4.6 we discuss the situation in which satisfaction of the constraint $\mathbf{x} \in C$ also has to be estimated, but for now we assume feasibility can be established without error. If the problem is more naturally formulated as a maximization, then stating it as $\min -\theta(\mathbf{x})$ converts it into the formulation (8.19).

Let \mathbf{x}^* denote the optimal scenario, which we initially assume is unique. To simplify notation let $\hat{\theta}(\mathbf{x}) = \hat{\theta}(\mathbf{x}; T, n, \mathbf{U})$ when we do not need to explicitly consider the other experiment design components T, n and \mathbf{U} .

How we attack (8.19) depends on a number of things, including the type of feasible region C (discrete or continuous, for instance), the number of feasible scenarios in C (finite, countably infinite, uncountably infinite), and what we know about $\theta(\mathbf{x})$ and $\hat{\theta}(\mathbf{x})$. We can usually expect that $\hat{\theta}(\mathbf{x})$ is an unbiased, or at least a consistent, estimator of $\theta(\mathbf{x})$. And the SO algorithm will simulate K scenarios $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K$ (where K may be random), and choose as the best

$$\hat{\mathbf{x}}^* = \operatorname{argmin}_{i=1,2,\dots,K} \hat{\theta}(\mathbf{x}_i).$$

In other words, the algorithm returns the scenario with the sample best performance among all of those it simulated, and the estimated objective function value of the selected scenario is $\hat{\theta}(\hat{\mathbf{x}}^*)$. The number of feasible scenarios simulated will always be finite, since we have to stop eventually. In some cases K may exhaust C , so that we simulate all of the feasible scenarios, and in other cases $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K$ is a subset of C that is generated by some sort of search.

As a specific example, consider a variation of the stochastic activity network (SAN) problem introduced in Sect. 3.4 and simulated in Sect. 4.4. The SO version of this problem is based on Henderson and Nelson (2006).

Recall that for the SAN

$$Y = \max\{A_1 + A_4, A_1 + A_3 + A_5, A_2 + A_5\}$$

is the total time to complete the project, where A_j is the time required for activity j (previously we denoted the activity times by X_j , but we use A_j here to avoid confusion with the scenario variable \mathbf{x}). Suppose that activity j is exponentially distributed with nominal mean τ_j , but we can reduce the mean to x_j at a cost of $c_j(\tau_j - x_j)$. In other words, c_j is the cost per unit reduction in the mean time of activity j . Thus, the project completion time as a function of the scenario $\mathbf{x} = (x_1, x_2, x_3, x_4, x_5)$ is

$$Y(\mathbf{x}) = \max\{A_1(x_1) + A_4(x_4), A_1(x_1) + A_3(x_3) + A_5(x_5), A_2(x_2) + A_5(x_5)\}$$

where $A_j(x_j)$ is exponentially distributed with mean x_j . Let $\theta(x_1, x_2, \dots, x_5) = E(Y(\mathbf{x}))$, the expected value of the time to complete the project. If we have a fixed total budget b for the project, then the SO problem is

$$\begin{aligned} \min \theta(x_1, x_2, \dots, x_5) & \quad (8.20) \\ \sum_{j=1}^5 c_j(\tau_j - x_j) & \leq b \\ x_j & \geq \ell_j, \quad j = 1, 2, 3, 4, 5 \end{aligned}$$

where ℓ_j is the smallest achievable mean activity time for activity j . Notice that in this formulation we can actually save cost by letting $x_j > \tau_j$ for some activities.

The natural objective-function estimator for scenario $\mathbf{x}_i = (x_{i1}, x_{i2}, \dots, x_{i5})$ is

$$\hat{\theta}(\mathbf{x}_i) = \frac{1}{n_i} \sum_{j=1}^{n_i} Y_j(\mathbf{x}_i)$$

where $Y_1(\mathbf{x}_i), Y_2(\mathbf{x}_i), \dots$ are i.i.d. replications of the time to complete the project for scenario \mathbf{x}_i .

The size of the budget b , and the way in which we reduce the mean activity time to x_i , lead to very different problems:

- If we reduce the mean of activity i by allocating extra workers to the activity, and our budget is so tight that we can only afford to add at most one worker to at

most one or two activities, then the number of feasible scenarios is small and we can likely simulate them all.

- If b is very large relative to the cost of an individual worker, and each worker assigned to activity i further reduces the mean time to complete the activity, then there might be so many feasible scenarios that we cannot simulate them all even though the number is finite. Therefore, the SO will require a search.
- If the reduction in mean activity time comes not from allocating workers, but from allocating some sort of capacity or power, then we might be able to treat x_i as continuous valued, implying that there might be uncountably many feasible scenarios.

No matter what algorithm we used to search for and simulate the collection of feasible scenarios $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K$, there are three fundamental types of errors that can occur in SO problems:

1. *The optimal scenario is never simulated*, meaning

$$\mathbf{x}^* \notin \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K\}.$$

This is also the case in deterministic nonlinear optimization problems when the feasible region cannot be exhausted, so we should not expect stochastic simulation optimization to be any easier.

2. *The best scenario that was simulated is not selected*, meaning

$$\widehat{\mathbf{x}}^* \neq \mathbf{x}_B = \operatorname{argmin}_{i=1,2,\dots,K} \theta(\mathbf{x}_i).$$

Because $\widehat{\theta}(\mathbf{x}_i)$ only estimates $\theta(\mathbf{x}_i)$, the best scenario that we simulated may not have the best *estimated* objective function value, and this causes us to select an inferior scenario. This error never occurs in deterministic optimization problems.

3. *The estimated objective function value of the selected scenario is not very accurate*, meaning

$$\left| \widehat{\theta}(\widehat{\mathbf{x}}^*) - \theta(\widehat{\mathbf{x}}^*) \right| \text{ is large.}$$

Because we select the scenario with the smallest *estimated* value, there is a natural bias toward scenarios whose estimated performance is better (smaller) than its true expected performance. Therefore, we tend to have an optimistic estimate of how good the selected scenario actually is; that is, we tend to have

$$\widehat{\theta}(\widehat{\mathbf{x}}^*) < \theta(\widehat{\mathbf{x}}^*).$$

Restating these errors in the context of the SAN optimization problem (8.20), they are:

1. The SO algorithm never discovers and simulates the best feasible setting of the mean activity times $\mathbf{x}^* = (x_1^*, x_2^*, \dots, x_5^*)$.
2. The best setting actually simulated, \mathbf{x}_B , does not have the smallest estimated mean project completion time among $\widehat{\theta}(\mathbf{x}_i), i = 1, 2, \dots, K$.

3. The sample mean activity time of the selected scenario $\widehat{\theta}(\widehat{\mathbf{x}}^*)$ is not very close to its true mean project completion time $\theta(\widehat{\mathbf{x}}^*)$.

Clearly we would like to control, measure or eliminate Errors 1–3. Guaranteed asymptotic convergence (see Sect. 8.3.1) is the primary way SO algorithms are designed to address Error 1. Correct-selection guarantees address Error 2, and may also address Error 3; see Sect. 8.3.2.

8.3.1 Convergence

Convergence guarantees say something about what would happen if the SO algorithm was allowed to run forever. “Run forever” often implies exploring more and more feasible scenarios and simulating them more and more thoroughly. We focus on two particular definitions of convergence; for a broader discussion of convergence see Andradóttir (2006a).

The typical SO algorithm is iterative, and on the r th iteration it reports an estimated optimal scenario $\widehat{\mathbf{x}}_r^*$, $r = 1, 2, \dots$. An iteration could consist of simulating new feasible scenarios, or refining the estimates of scenarios that have already been simulated, or both. Notice that SO algorithms can and often do revisit scenarios that have already been simulated to better estimate their objective function values, a feature that makes no sense for deterministic optimization. As a result, it need not be the case that $K = r$ (the number of distinct scenarios simulated need not be the same as the number of iterations of the algorithm), and typically $K \ll r$.

A very natural definition of convergence for SO is

$$\Pr \left\{ \lim_{r \rightarrow \infty} \widehat{\mathbf{x}}_r^* = \mathbf{x}^* \right\} = 1. \quad (8.21)$$

An algorithm with this property is *globally convergent* with probability 1. There exist globally convergent simulation optimization algorithms, particularly for the case when C contains a finite (even if very large) number of feasible scenarios. When the number of feasible scenarios is finite, convergence is usually proven by establishing that every feasible scenario will be simulated infinitely often in the limit. A good globally convergent algorithm assures global convergence even though it aggressively pursues improving scenarios. Although nothing like global convergence may actually happen in practice, it is reassuring to know that the algorithm would find the optimal scenario in the limit.

Next we consider *local convergence*. To do so we need the concept of a neighborhood of a feasible scenario \mathbf{x} , denoted by $N(\mathbf{x})$. The neighborhood of \mathbf{x} is a collection of feasible scenarios (so $N(\mathbf{x}) \subset C$) that are in some sense close to \mathbf{x} . For instance, if the components of \mathbf{x} are integers (such as the initial level of inventory of d products), then one definition of a neighborhood is scenarios that differ from \mathbf{x} by ± 1 in one component. That is, all scenarios of the form $\mathbf{x} \pm (0, 0, \dots, 0, 1, 0, \dots, 0)$ that are also feasible. A scenario \mathbf{x}' is *locally optimal* if

$$\theta(\mathbf{x}') \leq \theta(\mathbf{x}) \text{ for all } \mathbf{x} \in N(\mathbf{x}');$$

that is, \mathbf{x}' has as good or better (smaller) performance than the scenarios in its neighborhood. Of course, a scenario can be locally optimal but still not be very good overall.

Let $L \subset C$ be the set of locally optimal scenarios. Then local convergence is convergence to L . One definition is

$$\Pr\{\widehat{\mathbf{x}}_r^* \notin L \text{ infinitely often}\} = 0. \quad (8.22)$$

This means with probability 1 there is an iteration R such that for all iterations $r \geq R$ the estimated optimal scenario $\widehat{\mathbf{x}}_r^*$ is a locally optimal scenario. Local convergence is less satisfying than global convergence, but it also does not require that all feasible scenarios are simulated even in the limit.

While a convergence guarantee, particularly a global convergence guarantee, is clearly a desirable property, a convergent algorithm will still, in the end, simulate only $K < \infty$ scenarios and return the one with the best estimated value. Correct-selection guarantees, described in the next section, relate to what can be said about the K simulated scenarios. When combined with a locally convergent algorithm, correct selection can provide a statistical guarantee of optimality.

8.3.2 Correct Selection

“Correct selection” primarily addresses Error 2. Let $\mathbf{x}_B = \operatorname{argmin}_{i=1,2,\dots,K} \theta(\mathbf{x}_i)$ be the best scenario of those actually simulated by the SO algorithm; $\mathbf{x}_B = \mathbf{x}^*$ only if the algorithm actually simulated \mathbf{x}^* , and \mathbf{x}_B may not even be locally optimal. However, once the SO algorithm stops then selecting \mathbf{x}_B is the best that it can do.

Consider the correct-selection event

$$\begin{aligned} \text{CS} &= \{\text{select } \mathbf{x}_B\} \\ &= \{\widehat{\theta}(\mathbf{x}_B) < \widehat{\theta}(\mathbf{x}_i), i = 1, 2, \dots, K, i \neq B\} \\ &= \bigcap_{i=1, i \neq B}^K \{\widehat{\theta}(\mathbf{x}_B) < \widehat{\theta}(\mathbf{x}_i)\}. \end{aligned} \quad (8.23)$$

CS requires that when the optimization algorithm terminates the scenario with the best *estimated* performance is actually the best of the scenarios that were simulated. When this happens we will make a correct selection, and we would like a guarantee that $\Pr\{\text{CS}\}$ is high. Unfortunately, this sort of guarantee is affected by *multiplicity*, which is the effect that K has on the probability of an event like CS. To get some sense of this, consider two cases:

Independence: Suppose that the events $\{\widehat{\theta}(\mathbf{x}_B) < \widehat{\theta}(\mathbf{x}_i)\}$ are independent. Then

$$\Pr\{\text{CS}\} = \prod_{i=1, i \neq B}^K \Pr\{\widehat{\theta}(\mathbf{x}_B) < \widehat{\theta}(\mathbf{x}_i)\}.$$

Thus, the more feasible scenarios \mathbf{x}_i that we explore, the smaller the probability is that we select the correct one. For instance, if all of these events have probability $1 - \alpha$ then $\Pr\{\text{CS}\} = (1 - \alpha)^{K-1}$. This is unfortunate because exploring broadly increases the chance of uncovering the optimal scenario \mathbf{x}^* .

General but unknown dependence: Since the events $\{\widehat{\theta}(\mathbf{x}_B) < \widehat{\theta}(\mathbf{x}_i)\}$ all depend on the same $\widehat{\theta}(\mathbf{x}_B)$, it is unrealistic to think that they will be independent. When they are dependent, and we know nothing more than that, then the best we can say is that

$$\Pr\{\text{CS}\} \geq 1 - \sum_{i=1, i \neq B}^K \Pr\{\widehat{\theta}(\mathbf{x}_B) \geq \widehat{\theta}(\mathbf{x}_i)\}.$$

This is even more discouraging than the independent case since this lower bound, which is known as the Bonferroni inequality, can become negative (and therefore meaningless) if K is large.²

The curse of multiplicity is that *exploring broadly (large K) makes it difficult to select correctly* because the statistical errors accumulate. There are two direct ways to mitigate this problem and both attack the probability of making a mistake, $\Pr\{\widehat{\theta}(\mathbf{x}_B) \geq \widehat{\theta}(\mathbf{x}_i)\}$.

Suppose we can design our SO algorithm to guarantee that

$$\Pr\{\widehat{\theta}(\mathbf{x}_B) \geq \widehat{\theta}(\mathbf{x}_i)\} \leq \frac{\alpha}{K-1}$$

for all $\mathbf{x}_i \neq \mathbf{x}_B$. This means the chance of error in each comparison is very small. Then the Bonferroni inequality guarantees that $\Pr\{\text{CS}\} \geq 1 - (K-1)\alpha/(K-1) = 1 - \alpha$. For instance, if our simulation optimization explores $K = 1,001$ feasible scenarios and we want 95% confidence we have selected the best scenario of those simulated then we can achieve this if

$$\Pr\{\widehat{\theta}(\mathbf{x}_B) \geq \widehat{\theta}(\mathbf{x}_i)\} \leq \frac{0.05}{1000} = 0.00005$$

for each alternative scenario \mathbf{x}_i . Not surprisingly, obtaining such a small level of error typically requires a very large expenditure of simulation effort (replications or run length), so this approach tends to be viable only when K is small.

The second approach tries to decrease $\Pr\{\widehat{\theta}(\mathbf{x}_B) \geq \widehat{\theta}(\mathbf{x}_i)\}$ by altering the joint distribution of $\widehat{\theta}(\mathbf{x}_i; \mathbf{U}_i)$, $i = 1, 2, \dots, K$, where we now need to reintroduce the random numbers into the notation.

Suppose that we can represent the performance estimator as

$$\widehat{\theta}(\mathbf{x}_i; \mathbf{U}_i) = \theta(\mathbf{x}_i) + \varepsilon(\mathbf{U}_i)$$

² There are various expressions of the Bonferroni inequality, but the one we use here is $\Pr\{\bigcap_{i=1}^k \mathcal{E}_i\} \geq 1 - \sum_{i=1}^k \Pr\{\mathcal{E}_i^c\}$ where \mathcal{E}_i is an event and \mathcal{E}_i^c is its complement.

where $\varepsilon(\mathbf{U}_i)$ is a mean-zero random variable that accounts for the fact that the simulation output is stochastic. Notice that in this stylized representation the simulation noise ε does not depend on the scenario \mathbf{x}_i . Therefore, if we use the same pseudo-random numbers for each scenario, $\mathbf{U}_1 = \dots = \mathbf{U}_K = \mathbf{U}$, then

$$\begin{aligned} \Pr\left\{\widehat{\theta}(\mathbf{x}_B; \mathbf{U}) \geq \widehat{\theta}(\mathbf{x}_i; \mathbf{U})\right\} &= \Pr\{\theta(\mathbf{x}_B) + \varepsilon(\mathbf{U}) \geq \theta(\mathbf{x}_i) + \varepsilon(\mathbf{U})\} \\ &= \Pr\{\theta(\mathbf{x}_B) \geq \theta(\mathbf{x}_i)\} = 0. \end{aligned}$$

Thus, we never make a mistake no matter what K is!

We cannot expect the simulation noise to be independent of the scenario, or independent of the run length T_i or number of replications n_i for that matter. But this analysis does suggest that if the K scenarios behave similarly with respect to the random numbers, then we can decrease the probability of selection error by assigning the same random numbers to all scenarios. This technique is known as using *common random numbers* and we explore it more fully in Sect. 8.4.1.

How does the probability of correct selection relate to convergence? When a SO algorithm is globally convergent, but we cannot exhaust the feasible region C (as will always be the case when \mathbf{x} is continuous valued), then a correct-selection guarantee only provides inference with respect to the scenarios $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K$ that were actually simulated; it says nothing about overall optimality. For a locally convergent algorithm, a correct-selection guarantee can provide confidence that $\widehat{\mathbf{x}}^*$ is actually locally optimal provided all of its neighbors $N(\widehat{\mathbf{x}}^*)$ have been simulated. Since the number of scenarios in $N(\widehat{\mathbf{x}}^*)$ is typically much smaller than the number of scenarios in C , simulating all neighbors is often possible. However, the inference is regarding local optimality, not global. Only when the number of feasible scenarios is small enough that we can simulate every $\mathbf{x} \in C$ does correct selection relate directly to overall optimality.

SO algorithms that do not exhaust the feasible scenarios C typically do not provide a correct-selection guarantee when they terminate. However, an (often small) amount of additional simulation of $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K$ can be used to add this guarantee, and also control Error 3 which is that $\theta(\widehat{\mathbf{x}}^*)$ is estimated badly. This strategy is known as “cleaning up” after the SO algorithm, and we describe one specific way to do it later. See also Boesel et al. (2003).

8.4 Experiment Design for Simulation Optimization

Having established the objectives in SO, we now turn to experiment design. Recall that we estimate the objective function $\theta(\mathbf{x})$ by $\widehat{\theta}(\mathbf{x}_i) = \widehat{\theta}(\mathbf{x}_i; T_i, n_i, \mathbf{U}_i)$. Experiment design consists of choosing the scenarios to simulate $\mathbf{x}_i, i = 1, 2, \dots, K$, the allocation of simulation effort (n_i, T_i) and the assignment of random numbers \mathbf{U}_i to scenario \mathbf{x}_i . We begin with random-number assignment.

8.4.1 Random-Number Assignment

The discussion of correct selection in Sect. 8.3.2 showed why the assignment of random numbers to the simulation of alternative scenarios should be a part of experiment design. Here we delve more deeply into that idea by first connecting correct selection to variance reduction, and then describing how random-number assignment reduces variance.

We will assume two blocks of random numbers, say \mathbf{U}_i and \mathbf{U}_h , are independent if $i \neq h$. As a practical matter we know that \mathbf{U}_i is actually a block of deterministic pseudorandom numbers $\mathbf{u}_i = (u_{i1}, u_{i2}, \dots, u_{is})$. We assume that s is so large that if we assign \mathbf{u}_i to some purpose we will never use all s pseudorandom numbers in the block. We treat the random numbers \mathbf{U}_i and \mathbf{U}_h as independent because they actually correspond to nonoverlapping sequences of pseudorandom numbers \mathbf{u}_i and \mathbf{u}_h . This allows us to “assign” either common or distinct blocks of random numbers to scenarios even though “assigning” anything that is truly random is a contradiction. The assignment is done by specifying seeds or streams as described in Sect. 6.5.3.

In Sect. 8.3.2 we focussed on $\Pr \left\{ \hat{\theta}(\mathbf{x}_B) \geq \hat{\theta}(\mathbf{x}_i) \right\}$, the probability that the estimated objective function value of the best scenario we simulated is larger than the estimated objective function value of a scenario that is actually inferior. Since we are minimizing, and since we select the scenario that appears to be the best, this event corresponds to an incorrect selection.

Suppose that these estimators are averages of a large number of more basic observations (e.g., replications) so that we are willing to treat them as normally distributed due to the central limit theorem. Then

$$\begin{aligned}
 \Pr \left\{ \hat{\theta}(\mathbf{x}_B) \geq \hat{\theta}(\mathbf{x}_i) \right\} &= \Pr \left\{ \hat{\theta}(\mathbf{x}_B) - \hat{\theta}(\mathbf{x}_i) \geq 0 \right\} \\
 &= \Pr \left\{ \frac{\hat{\theta}(\mathbf{x}_B) - \hat{\theta}(\mathbf{x}_i) - (\theta(\mathbf{x}_B) - \theta(\mathbf{x}_i))}{\sqrt{\text{Var} \left[\hat{\theta}(\mathbf{x}_B) - \hat{\theta}(\mathbf{x}_i) \right]}} \geq \frac{-(\theta(\mathbf{x}_B) - \theta(\mathbf{x}_i))}{\sqrt{\text{Var} \left[\hat{\theta}(\mathbf{x}_B) - \hat{\theta}(\mathbf{x}_i) \right]}} \right\} \\
 &= \Pr \left\{ Z \geq \frac{\theta(\mathbf{x}_i) - \theta(\mathbf{x}_B)}{\sqrt{\text{Var} \left[\hat{\theta}(\mathbf{x}_B) - \hat{\theta}(\mathbf{x}_i) \right]}} \right\} \tag{8.24}
 \end{aligned}$$

where Z is a $N(0, 1)$ random variable. Since $\theta(\mathbf{x}_i) - \theta(\mathbf{x}_B) > 0$, Expression (8.24) is an increasing function of $\text{Var} \left[\hat{\theta}(\mathbf{x}_B) - \hat{\theta}(\mathbf{x}_i) \right]$. Stated differently, the smaller the $\text{Var} \left[\hat{\theta}(\mathbf{x}_B) - \hat{\theta}(\mathbf{x}_i) \right]$ is, the smaller the probability of an incorrect selection (or the greater the chance of a correct selection). This makes intuitive sense: the better our estimate of the difference $\hat{\theta}(\mathbf{x}_B) - \hat{\theta}(\mathbf{x}_i)$ is, as measured by variance, the more likely we are to choose the one that is actually better.

The variance of the difference is³

$$\text{Var} \left[\widehat{\theta}(\mathbf{x}_B) - \widehat{\theta}(\mathbf{x}_i) \right] = \text{Var} \left[\widehat{\theta}(\mathbf{x}_B) \right] + \text{Var} \left[\widehat{\theta}(\mathbf{x}_i) \right] - 2\text{Cov} \left[\widehat{\theta}(\mathbf{x}_B), \widehat{\theta}(\mathbf{x}_i) \right]. \quad (8.25)$$

If each scenario \mathbf{x}_i is assigned independent random numbers \mathbf{U}_i , then the estimators are independent and $\text{Cov} \left[\widehat{\theta}(\mathbf{x}_i; \mathbf{U}_i), \widehat{\theta}(\mathbf{x}_B; \mathbf{U}_h) \right] = 0$. Common random numbers means letting $\mathbf{U}_i = \mathbf{U}_B = \mathbf{U}$. The desired effect so as to increase the probability of correct selection is $\text{Cov} \left[\widehat{\theta}(\mathbf{x}_i, \mathbf{U}), \widehat{\theta}(\mathbf{x}_B, \mathbf{U}) \right] > 0$ which reduces the variance of the difference and increases the probability of correct selection. This is the connection between variance reduction and correct selection.

To gain some intuition, consider the output from the SAN simulation, rewritten in terms of the random numbers $(U_1, U_2, U_3, U_4, U_5)$ needed to generate the exponentially distributed activity times:

$$\begin{aligned} Y(\mathbf{x}) &= \max \{ A_1(x_1) + A_4(x_4), A_1(x_1) + A_3(x_3) + A_5(x_5), A_2(x_2) + A_5(x_5) \} \\ &= \max \{ -\ln(1 - U_1)x_1 - \ln(1 - U_4)x_4, \\ &\quad -\ln(1 - U_1)x_1 - \ln(1 - U_3)x_3 - \ln(1 - U_5)x_5, \\ &\quad -\ln(1 - U_2)x_2 - \ln(1 - U_5)x_5 \} \end{aligned}$$

where x_j is the mean time for activity j . The use of common random numbers means that the simulations of scenarios \mathbf{x}_i and \mathbf{x}_h differ only in the values of x that multiply $-\ln(1 - U)$. Thus, a collection of random numbers $(U_1, U_2, U_3, U_4, U_5)$ that make $Y(\mathbf{x}_i)$ larger than its expected value will tend to make $Y(\mathbf{x}_h)$ larger than its expected value also. This effect is particularly easy to see if for \mathbf{x}_i we have $x_{i1} = x_{i2} = \dots = x_{i5} = x_i$, and for \mathbf{x}_h we have $x_{h1} = x_{h2} = \dots = x_{h5} = x_h$. Then

$$\begin{aligned} Y(\mathbf{x}_h) &= \max \{ -\ln(1 - U_1)x_h - \ln(1 - U_4)x_h, \\ &\quad -\ln(1 - U_1)x_h - \ln(1 - U_3)x_h - \ln(1 - U_5)x_h, \\ &\quad -\ln(1 - U_2)x_h - \ln(1 - U_5)x_h \} \\ &= \frac{x_h}{x_i} \max \{ -\ln(1 - U_1)x_i - \ln(1 - U_4)x_i, \\ &\quad -\ln(1 - U_1)x_i - \ln(1 - U_3)x_i - \ln(1 - U_5)x_i, \\ &\quad -\ln(1 - U_2)x_i - \ln(1 - U_5)x_i \} \\ &= \frac{x_h}{x_i} Y(\mathbf{x}_i). \end{aligned}$$

Therefore,

$$\text{Cov}[Y(\mathbf{x}_i), Y(\mathbf{x}_h)] = \frac{x_h}{x_i} \text{Var}[Y(\mathbf{x}_i)] > 0.$$

³ A basic result from mathematical statistics is that for random variables A and B , $\text{Var}(A \pm B) = \text{Var}(A) + \text{Var}(B) \pm 2\text{Cov}(A, B)$.

Although it is more difficult to show mathematically, the covariance is still positive (although perhaps not as large) even when \mathbf{x}_i and \mathbf{x}_h do not have this special structure.

Two aspects of this example turn out to be important more generally: Notice that the output $Y(\mathbf{x})$ is monotonic in each of the random numbers $(U_1, U_2, U_3, U_4, U_5)$, and the same random number was used for the same purpose when simulating scenarios \mathbf{x}_i and \mathbf{x}_h ; e.g., we used U_j to generate activity time A_j in each scenario. Monotonicity is the result of two features of the simulation: (1) The inverse-cdf method $A_j = F^{-1}(U_j)$ was used to generate the inputs A_j and it is always nondecreasing in U_j , and (2) the structure of the simulation itself caused the output to be monotonic in each of the inputs; that is, $Y = \max\{A_1 + A_4, A_1 + A_3 + A_5, A_2 + A_5\}$ is nondecreasing in each A_i . Monotonicity is central to inducing positive covariance, and synchronization (using the same random number for the same purpose in each scenario) tends to maximize the effect. This is shown formally in Glasserman and Yao (1992), but also makes intuitive sense because positive covariance means that the two quantities tend to go up and down together.

Synchronization is easy for the SAN simulation because each replication of each scenario uses exactly five random numbers. Thus, if the same starting seed or stream is assigned to each scenario, then the simulation of each scenario starts with the random number (say) U_1 , and the j th replication uses the random numbers $U_{(j-1)5+1}, U_{(j-1)5+2}, U_{(j-1)5+3}, U_{(j-1)5+4}, U_{(j-1)5+5}$ to generate activity times A_1, A_2, A_3, A_4, A_5 , respectively, regardless the value of \mathbf{x} .

Unfortunately, for other simulations synchronization is not as easy. To illustrate this, consider simulating two queueing systems that have the following features to compare their mean waiting times:

1. They have identical renewal arrival processes.
2. Scenario $x = 1$ has one fast server, while scenario $x = 2$ has two slower ones. Their service-time distributions are from the same family, but with different parameters.
3. Customers are not treated equally; instead, customers who are older are moved to the head of the queue. We have an input distribution for customer age.
4. We will simulate $n = 100$ replications of each scenario, and each replication ends after $T = 8$ h of simulated time.

Notice that the simulation does have some structural monotonicity: longer interarrival times lead to shorter waiting times (monotone decreasing); and longer service times lead to longer waiting times (monotone increasing). However, assigning the same starting seed or stream to each scenario is not sufficient here. Table 8.2 shows what could happen if we do so and just consume the random numbers (RNs) as needed by the simulation. Because the number of servers differ, rather quickly RNs are being used for different purposes. For instance, U_7 is used to generate a customer age in scenario 1 and an interarrival-time in scenario 2. Therefore, the structural monotonicity of the queueing system is not fully exploited. Further, since the first replication of scenario 1 consumes fewer random numbers than scenario 2 (because fewer customers are able to be served in $T = 8$ h), the synchronization

Table 8.2 A possible sequence of events and the random numbers consumed by the first and second replications of scenarios 1 and 2 using a single random number stream

Scenario 1 (one server)			Scenario 2 (two servers)		
Event	RN	Purpose	Event	RN	Purpose
Arrival	U_1	Interarrival	Arrival	U_1	Interarrival
	U_2	Age		U_2	Age
	U_3	Service time		U_3	Service time
Arrival	U_4	Interarrival	Arrival	U_4	Interarrival
	U_5	Age		U_5	Age
Arrival	U_6	Interarrival	Arrival	U_6	Service time
	U_7	Age		U_7	Interarrival
End service	U_8	Service time	End service	U_8	Age
	U_9	Service time		U_9	Service time
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
End service	U_{123}	Service time	End service	U_{139}	Service time
End of replication 1			End of replication 1		
Arrival	U_{124}	Interarrival	Arrival	U_{140}	Interarrival
	U_{125}	Age		U_{141}	Age
	U_{126}	Service time		U_{142}	Service time
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots

is even worse in the second replication. As we make more and more replications there will tend to be less and less synchronization, and smaller and smaller positive covariance.

The reason for pseudorandom-number generators having multiple starting seeds or streams is to facilitate synchronization. By assigning a distinct stream to each input process (arrivals, age and service time) we guarantee that each pseudorandom number is used for the same purpose in each scenario. To maintain this synchronization across replications, we also need distinct starting streams or seeds for each replication. Thus, implementing common random numbers for the queueing example requires $100 \times 3 = 300$ starting seeds or streams for full synchronization, which is easily doable with a pseudorandom-number generator like MRG32k3a in L'Ecuyer (1999). See L'Ecuyer et al. (2002) for an efficient implementation of multiple streams.

Partial synchronization is better than no synchronization at all, but the benefits of full synchronization can be substantial. Shechter et al. (2006) describe a simulation that compares survival rates of an HIV cohort under two different treatment policies using independent simulations (different random numbers assigned to policies 1 and 2), and using common random numbers with partial synchronization and full synchronization. Relative to independent simulations, using common random numbers with full synchronization reduced the variance of the difference estimate by 93%, but only by 37% with partial synchronization. See Kelton (2006) for further discussion about synchronization strategies.

Clearly, variate generation via the inverse cdf also aids synchronization since each input variate is a function of one U . However, it is not always essential. In

the queueing example the arrival process input model is the same for scenarios 1 and 2. Thus, as long as the arrival process has a separate random number stream, each scenario will experience exactly the same sequence of arrivals no matter what method is used to generate the interarrival times.

More care is required for the service times because the distribution changes. When the distribution changes but remains in the same family it is sometimes possible to view the inputs for both scenarios as a transformation of a common base random variable. Examples include the normal distribution with mean μ and variance σ^2 since if Z is $N(0, 1)$ then

$$\mu + \sigma Z$$

is $N(\mu, \sigma^2)$. Thus, as long as we generate the base Z using the same method and the same stream in both scenarios 1 and 2, we achieve full synchronization (and strong positive correlation). The same reasoning can be applied to lognormal variates, since they are transformations of normals.

Even when there is no simple expression for the inverse cdf, approximations and numerical inversion are possible (see Sect. 6.4). This will likely be slower than, say, a rejection method, but the additional effort may be worthwhile if a significant variance reduction can be achieved. When inversion is simply not possible, then Kachitvichyanukul and Schmeiser (1990) show that there are ways to create rejection-type algorithms that still maintain synchronization and monotonicity.

While our focus here has been on the effect of common random numbers on the probability of correct selection, the variance reduction is also relevant when the goal is simply to estimate the difference in performance between two or more scenarios. Suppose that we have just two scenarios, \mathbf{x}_1 and \mathbf{x}_2 , for the mean activity times in the SAN simulation, and we simulate n replications of each: $Y_1(\mathbf{x}_1), Y_2(\mathbf{x}_1), \dots, Y_n(\mathbf{x}_1)$ and $Y_1(\mathbf{x}_2), Y_2(\mathbf{x}_2), \dots, Y_n(\mathbf{x}_2)$. Let $D_j = Y_j(\mathbf{x}_1) - Y_j(\mathbf{x}_2)$, $j = 1, 2, \dots, n$ be the differences between corresponding replications of the two scenarios. Then the paired- t confidence interval for the difference $\theta(\mathbf{x}_1) - \theta(\mathbf{x}_2)$ is

$$\bar{D} \pm t_{1-\alpha/2, n-1} \frac{S_D}{\sqrt{n}} \quad (8.26)$$

where

$$\bar{D} = \frac{1}{n} \sum_{j=1}^n D_j$$

and

$$S_D^2 = \frac{1}{n-1} \sum_{j=1}^n (D_j^2 - \bar{D})^2.$$

This CI is valid with or without common random numbers, provided that the D_j are approximately normally distributed or n is large. The effect of common random numbers shows up in the $\text{Var}[D_j] = \text{Var}[Y_j(\mathbf{x}_1) - Y_j(\mathbf{x}_2)]$, which is estimated by S_D^2 .

Thus, we expect S_D^2 to be smaller, making the CI shorter, when common random numbers are used, and this makes differences easier to detect.

The pairing of replications in (8.26) is important under common random numbers. If we have synchronized correctly then $Y_j(\mathbf{x}_1)$ and $Y_j(\mathbf{x}_2)$ are dependent and positively correlated, while $Y_j(\mathbf{x}_1)$ and $Y_h(\mathbf{x}_2)$ are independent for different replications $j \neq h$. Therefore, we can treat the D_1, D_2, \dots, D_n as i.i.d.

For example, for the SAN we simulated two scenarios $\mathbf{x}_1 = (1, 1, 1, 1, 1)$ and $\mathbf{x}_2 = (0.5, 0.5, 1, 1.2, 1.2)$ with and without common random numbers. Based on 100 replications the estimated correlation between $Y_j(\mathbf{x}_1)$ and $Y_j(\mathbf{x}_2)$ was 0.93, reducing the half width of the confidence interval for the difference from ± 0.46 to ± 0.12 when the estimated difference is about $\bar{D} = 0.30$.

8.4.2 Design and Analysis for Correct Selection

Suppose that there are K scenarios $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K$ and we would like a guarantee of selecting the best among them. These scenarios arise either because there are exactly K feasible scenarios and all of them will be simulated, or because they are the K scenarios in C that were actually simulated during some sort of search. In the SAN optimization these would be K settings for the mean activity times that can be attained within the available budget.

The formal proofs that the procedures presented in this section provide their stated guarantees assume that the output data from each scenario \mathbf{x}_i , $Y_1(\mathbf{x}_i), Y_2(\mathbf{x}_i), \dots, Y_{n_i}(\mathbf{x}_i)$, are i.i.d. normally distributed with mean $\theta(\mathbf{x}_i)$ and finite variance. Independence is assured if they are the results from different replications, and may be approximately correct if they are batch statistics from a single replication of a steady-state simulation (see Sect. 8.2). Further, in the procedures that follow $\theta(\mathbf{x}_i)$ is estimated by the sample mean

$$\hat{\theta}(\mathbf{x}_i) = \bar{Y}(\mathbf{x}_i; n_i) = \frac{1}{n_i} \sum_{j=1}^{n_i} Y_j(\mathbf{x}_i).$$

As in Sect. 8.3.2, we let \mathbf{x}_B denote the unknown best scenario; that is,

$$\theta(\mathbf{x}_B) = \min_{i=1,2,\dots,K} \theta(\mathbf{x}_i).$$

There are numerous procedures available for many variations of this type of problem, and asymptotic justifications for cases in which normality and independence do not apply; some general references are Bechhofer et al. (1995); Frazier (2010); Goldsman and Nelson (1998) and Kim and Nelson (2006). We present just two procedures that have been useful in practice, and that illustrate some key ideas.

8.4.2.1 Subset Selection

A subset selection procedure delivers a set of feasible scenarios $I \subseteq \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K\}$ with a guarantee that

$$\Pr\{\mathbf{x}_B \in I\} \geq 1 - \alpha.$$

In the best case I contains a single scenario, which is \mathbf{x}_B with probability at least $1 - \alpha$. More typically, subset selection eliminates or screens out scenarios that (with high statistical confidence) are not the best so that attention can be focused on a much smaller group. The specific procedure below is particularly useful when $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K$ are the result of a search of C because it can allow the scenarios to have different numbers of observations.

1. Given $n_i \geq 2$ observations from scenario \mathbf{x}_i , set

$$t_i = t_{(1-\alpha)^{\frac{1}{K-1}}, n_i-1}$$

the $(1 - \alpha)^{\frac{1}{K-1}}$ quantile of the t distribution with $n_i - 1$ degrees of freedom, for $i = 1, 2, \dots, K$.

2. Calculate the sample means $\bar{Y}(\mathbf{x}_i; n_i)$ and sample variances

$$S^2(\mathbf{x}_i) = \frac{1}{n_i - 1} \sum_{j=1}^{n_i} (Y_j(\mathbf{x}_i) - \bar{Y}(\mathbf{x}_i; n_i))^2$$

for $i = 1, 2, \dots, K$, and also the threshold

$$W_{ih} = \left(t_i^2 \frac{S^2(\mathbf{x}_i)}{n_i} + t_h^2 \frac{S^2(\mathbf{x}_h)}{n_h} \right)^{1/2}$$

for all $i \neq h$.

3. Form the subset

$$I = \{ \mathbf{x}_i : \bar{Y}(\mathbf{x}_i; n_i) \leq \bar{Y}(\mathbf{x}_h; n_h) + W_{ih} \text{ for all } h \neq i \}.$$

The rule for which scenarios are included in the subset is simple: Include \mathbf{x}_i if

$$\bar{Y}(\mathbf{x}_i; n_i) \leq \bar{Y}(\mathbf{x}_h; n_h) + W_{ih}, h \neq i.$$

That is, we retain scenario \mathbf{x}_i if its sample mean is smaller than all of the other sample means adjusted with a positive quantity W_{ih} that accounts for estimation error. Why does this work? The following argument is behind many subset selection procedures:

$$\begin{aligned} & \Pr\{\mathbf{x}_B \in I\} \\ &= \Pr\{ \bar{Y}(\mathbf{x}_B; n_B) \leq \bar{Y}(\mathbf{x}_h; n_h) + W_{Bh}, h \neq B \} \\ &= \Pr\{ \bar{Y}(\mathbf{x}_B; n_B) - \bar{Y}(\mathbf{x}_h; n_h) - [\theta(\mathbf{x}_B) - \theta(\mathbf{x}_h)] \leq W_{Bh} - [\theta(\mathbf{x}_B) - \theta(\mathbf{x}_h)], h \neq B \} \\ &\geq \Pr\{ \bar{Y}(\mathbf{x}_B; n_B) - \bar{Y}(\mathbf{x}_h; n_h) - [\theta(\mathbf{x}_B) - \theta(\mathbf{x}_h)] \leq W_{Bh}, h \neq B \}. \end{aligned} \quad (8.27)$$

The Inequality (8.27) follows because $-\left[\theta(\mathbf{x}_B) - \theta(\mathbf{x}_h)\right] \geq 0$, so removing it from the right-hand side makes the event more difficult to satisfy. The statistic

$$\bar{Y}(\mathbf{x}_i; n_i) - \bar{Y}(\mathbf{x}_h; n_h) - [\theta(\mathbf{x}_i) - \theta(\mathbf{x}_h)]$$

has mean 0 for all $i \neq h$, allowing the W_{ih} 's to be derived that give the desired probability based only on their variances.

The formal proof of validity is in Boesel et al. (2003), and it assumes that common random numbers are *not* used in the experiment design. This makes sense because the number of observations from each scenario can be different so the replications from different systems cannot be paired. If the number of observations are forced to be equal to n_0 , say, and common random numbers are employed, then the procedure is valid if W_{ih} is replaced with

$$W'_{ih} = \left(t^2 \frac{S_{ih}^2}{n_0} \right)^{1/2}$$

with

$$t = t_{1-\alpha/(K-1), n_0-1}$$

and

$$S_{ih}^2 = \frac{1}{n_0 - 1} \sum_{j=1}^{n_0} (Y_j(\mathbf{x}_i) - Y_j(\mathbf{x}_h) - [\bar{Y}(\mathbf{x}_i; n_0) - \bar{Y}(\mathbf{x}_h; n_0)])^2.$$

The effect of common random numbers is to reduce the size of the selected subset since it reduces the variance of the difference which in turn reduces W_{ih} . See Nelson et al. (2001).

8.4.2.2 Selection of the Best

Subset selection is an analysis method that takes whatever output data are available on the scenarios $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K$ and provides inference on which ones are not the best. In this section we describe a design and analysis procedure that controls the amount of simulation effort expended and delivers a single selection at the end. To make this happen the user is required to specify the smallest difference in expected performance that is practically important, which we will call δ . Here are some examples:

In a queueing simulation where waiting times are on the order of tens of minutes, a difference of less than 1/2 min (30 s) in mean waiting time might not be practically important, but a difference of more than that matters; therefore, we might set $\delta = 1/2$ min. In a reliability simulation where time-to-failure is on the order of weeks, differences between scenarios of 12 h or more in mean time-to-failure might be meaningful to the system designers; thus, we could set $\delta = 12$ h. And in a project planning optimization if the project takes a year or more to complete, then we might be satisfied with a scenario \mathbf{x}_i that is not optimal if its mean time-to-completion

$\theta(\mathbf{x}_i)$ is no more than $\delta = 7$ days longer than the optimal mean time $\theta(\mathbf{x}_B)$. The choice of δ is entirely up to the user, and depends on the situation.

Each iteration of the procedure below takes a single observation from scenarios that have not yet been screened out, applies a type of subset selection, and continues until there is only one scenario in the subset. It is valid with or without common random numbers, and can be very efficient (i.e., terminate quickly) if common random numbers are employed. The guarantee provided is that

$$\Pr\{\text{select } \mathbf{x}_B | \theta(\mathbf{x}_i) - \theta(\mathbf{x}_B) \geq \delta, \forall i \neq B\} \geq 1 - \alpha.$$

If $\theta(\mathbf{x}_i) - \theta(\mathbf{x}_B) < \delta$ for some i , then it will select one of the scenarios within δ of the best. The specification of a practically significant difference δ is what allows the procedure to guarantee to terminate in finite time. Such procedures are called *indifference-zone selection procedures* and δ is what defines the indifference (not practically important) zone.

1. Specify a common first-stage number of observations $n_0 \geq 2$. Set

$$\eta = \frac{1}{2} \left[\left(\frac{2\alpha}{K-1} \right)^{-2/(n_0-1)} - 1 \right].$$

2. Let $I = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K\}$ be the set of scenarios still in contention, and let $t^2 = 2\eta(n_0 - 1)$. Obtain n_0 observations $Y_j(\mathbf{x}_i), j = 1, 2, \dots, n_0$ from each scenario $\mathbf{x}_i \in I$ and compute $\bar{Y}(\mathbf{x}_i; n_0)$. For all $i \neq h$ calculate

$$S_{ih}^2 = \frac{1}{n_0 - 1} \sum_{j=1}^{n_0} (Y_j(\mathbf{x}_i) - Y_j(\mathbf{x}_h) - [\bar{Y}(\mathbf{x}_i; n_0) - \bar{Y}(\mathbf{x}_h; n_0)])^2,$$

the sample variance of the difference between scenarios i and h . Set $r = n_0$.

3. Set $I^{\text{old}} = I$. Let

$$I = \left\{ \mathbf{x}_i : \mathbf{x}_i \in I^{\text{old}} \text{ and } \bar{Y}(\mathbf{x}_i; r) \leq \bar{Y}(\mathbf{x}_h; r) + W_{ih}(r), \forall h \in I^{\text{old}}, h \neq i \right\},$$

where

$$W_{ih}(r) = \max \left\{ 0, \frac{\delta}{2r} \left(\frac{t^2 S_{ih}^2}{\delta^2} - r \right) \right\}.$$

4. If $|I| = 1$, then stop and select the scenario in I as the best. Otherwise, take one additional observation $Y_{r+1}(\mathbf{x}_i)$ from each scenario $\mathbf{x}_i \in I$, update the sample means, set $r = r + 1$ and go to Step 3.

This procedure, due to Kim and Nelson (2001), applies subset selection iteratively as new observations are obtained, with the threshold $W_{ih}(r)$ shrinking to 0 so that eventually the scenario with the smallest sample mean is selected. Why does this work?

Recall the correct selection event

$$\begin{aligned} \text{CS} &= \{\text{select } \mathbf{x}_B\} \\ &= \left\{ \widehat{\theta}(\mathbf{x}_B) < \widehat{\theta}(\mathbf{x}_i), i = 1, 2, \dots, K, i \neq B \right\} \\ &= \bigcap_{i=1, i \neq B}^K \left\{ \widehat{\theta}(\mathbf{x}_B) < \widehat{\theta}(\mathbf{x}_i) \right\} \end{aligned}$$

and consider the

$$\Pr \left\{ \widehat{\theta}(\mathbf{x}_B) < \widehat{\theta}(\mathbf{x}_i) \right\}.$$

Since we operate under the assumption that $\theta(\mathbf{x}_i) - \theta(\mathbf{x}_B) \geq \delta$, we have

$$\begin{aligned} &\Pr \left\{ \widehat{\theta}(\mathbf{x}_B) < \widehat{\theta}(\mathbf{x}_i) \right\} \\ &= \Pr \left\{ \widehat{\theta}(\mathbf{x}_B) - \widehat{\theta}(\mathbf{x}_i) < 0 \right\} \\ &= \Pr \left\{ \widehat{\theta}(\mathbf{x}_B) - \widehat{\theta}(\mathbf{x}_i) - [\theta(\mathbf{x}_B) - \theta(\mathbf{x}_i)] < -[\theta(\mathbf{x}_B) - \theta(\mathbf{x}_i)] \right\} \\ &\geq \Pr \left\{ \widehat{\theta}(\mathbf{x}_B) - \widehat{\theta}(\mathbf{x}_i) - [\theta(\mathbf{x}_B) - \theta(\mathbf{x}_i)] \leq \delta \right\} \end{aligned} \quad (8.28)$$

where the Inequality (8.28) follows because we assume that the minimum difference is $\geq \delta$. As in subset selection, the statistic

$$\widehat{\theta}(\mathbf{x}_B) - \widehat{\theta}(\mathbf{x}_i) - (\theta(\mathbf{x}_B) - \theta(\mathbf{x}_i))$$

has mean 0, and since δ is known the procedure can be designed to provide the desired guarantee by considering only δ , the variances and the number of replications.

8.4.2.3 Example

We illustrate these procedures with the SAN optimization, supposing that the project budget allows reducing the mean activity time for only one activity; the feasible scenarios are shown below:

	x_1	x_2	x_3	x_4	x_5
x_{i1}	0.5	1	1	1	1
x_{i2}	1	0.5	1	1	1
x_{i3}	1	1	0.75	1	1
x_{i4}	1	1	1	0.5	1
x_{i5}	1	1	1	1	0.5

For subset selection, we independently simulated all $K = 5$ scenarios for $n = 100$ replications, giving the following summary statistics:

	\mathbf{x}_1	\mathbf{x}_2	\mathbf{x}_3	\mathbf{x}_4	\mathbf{x}_5
$\bar{Y}(\mathbf{x}_i, n_i)$	3.045384364	3.556953882	3.313437416	3.361689108	2.772924336
$S^2(\mathbf{x}_i)$	2.318117804	2.627223846	2.946949109	2.915804886	2.099817647

Clearly, \mathbf{x}_5 has the smallest sample mean, but which others are close enough that they cannot be eliminated from being the best? For confidence level $1 - \alpha = 0.95$, the subset-selection procedure designed for use without common random numbers gives the selected subset as $I = \{\mathbf{x}_1, \mathbf{x}_5\}$. Why is, say, \mathbf{x}_3 eliminated but not \mathbf{x}_1 ?

Notice that $(1 - \alpha)^{1/(K-1)} = 0.95^{1/4} = 0.987$, so the critical value for each scenario is $t_i^2 = t_{0.987,99}^2 = 5.145$. Scenario \mathbf{x}_3 is eliminated because

$$\begin{aligned} \bar{Y}(\mathbf{x}_3, n_3) &\approx 3.313 \not\leq \bar{Y}(\mathbf{x}_5, n_5) + \left(t_3^2 \frac{S^2(\mathbf{x}_3)}{n_3} + t_5^2 \frac{S^2(\mathbf{x}_5)}{n_5} \right)^{1/2} \\ &\approx 2.773 + \left(5.145 \frac{2.947}{100} + 5.145 \frac{2.100}{100} \right)^{1/2} \approx 3.282. \end{aligned}$$

However, \mathbf{x}_1 cannot be eliminated because

$$\bar{Y}(\mathbf{x}_1, n_1) \approx 3.045 \leq \bar{Y}(\mathbf{x}_5, n_5) + \left(t_1^2 \frac{S^2(\mathbf{x}_1)}{n_1} + t_5^2 \frac{S^2(\mathbf{x}_5)}{n_5} \right)^{1/2} \approx 3.250.$$

The guarantee is that the best scenario \mathbf{x}_B is either \mathbf{x}_1 or \mathbf{x}_5 with 95% confidence. In fact, it is easy to see that $\theta(\mathbf{x}_1) = \theta(\mathbf{x}_5)$, and a bit more difficult to see that they are the best scenarios, so this is a correct selection.

Next we applied selection of the best at confidence level $1 - \alpha = 0.95$, $n_0 = 10$ initial replications from each scenario, practically significant difference $\delta = 0.1$, and using common random numbers. The relevant constants are

$$\eta = \frac{1}{2} \left[\left(\frac{2\alpha}{K-1} \right)^{-2/(n_0-1)} - 1 \right] = \frac{1}{2} \left[\left(\frac{2(0.05)}{4} \right)^{-2/9} - 1 \right] \approx 0.635$$

and $t^2 = 2\eta(n_0 - 1) \approx 11.429$. At each iteration the procedure compares the sample means of all scenarios to each other. Comparison between, say, scenarios \mathbf{x}_1 and \mathbf{x}_5 requires the sample variance of the difference $S_{15}^2 \approx 0.294$, giving the decreasing threshold

$$\begin{aligned} W_{15}(r) &= \max \left\{ 0, \frac{\delta}{2r} \left(\frac{t^2 S_{15}^2}{\delta^2} - r \right) \right\} \\ &= \max \left\{ 0, \frac{0.1}{2r} (336.2 - r) \right\}. \end{aligned}$$

The thresholds $W_{12}(r)$, $W_{13}(r)$ and $W_{14}(r)$ are computed in a similar fashion. We eliminate scenario \mathbf{x}_1 at iteration r if

$$\bar{Y}(\mathbf{x}_1; r) \not\leq \bar{Y}(\mathbf{x}_h; r) + W_{1h}(r)$$

for some other scenario $h \neq 1$.

The procedure selected \mathbf{x}_5 as the best after 171, 10, 33, 109, 171 replications from $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5$, respectively. This means that \mathbf{x}_2 was eliminated after the initial sample, and then (in order) $\mathbf{x}_3, \mathbf{x}_4$ and \mathbf{x}_1 were eliminated. Why was scenario \mathbf{x}_1 eliminated if it is equivalent to \mathbf{x}_5 ? Because after 171 replications the procedure was able (with 95% confidence) to decide that \mathbf{x}_5 is either the best or within $\delta = 0.1$ of the best, and that is all that is required to make a correct selection.

8.4.2.4 Bayesian Procedures

There are two basic paradigms in designing for correct selection: frequentist (described above) and Bayesian. A Bayesian procedure consists of a sequence of decisions; the decisions include which scenario \mathbf{x} to simulate next, and possibly whether or not to stop the procedure and select a scenario. The key to making good decisions is that in the Bayesian formulation uncertainty about $\theta(\mathbf{x})$ is represented as a prior probability distribution on its value, which is updated using Bayes rule to a posterior distribution as simulation observations are obtained. A good starting point for Bayesian selection of the best is Frazier (2010).

An advantage of the Bayesian formulation is its flexibility; many kinds of information or knowledge about the problem can be incorporated into the prior beliefs, leading to substantial gains in efficiency. In addition, the procedure can be driven by different objectives: maximizing the posterior probability of correct selection and minimizing the opportunity cost for the selected scenario under a given simulation budget are two examples.

No procedure, frequentist or Bayesian, can dominate in all situations. Bayesian procedures are often more efficient in terms of the total number of observations required to make a selection. However, they do not provide the type of correct-selection guarantee that the frequentist procedures provide.

8.4.2.5 Cleaning up after Simulation Optimization

The procedures described in this section can be used for simulation optimization when $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K\}$ are the only feasible scenarios and all of them will be simulated. When this is the case, the three types of the SO error—failure to simulate the optimal scenario, failure to select the best scenario that was simulated, and poorly estimating the performance of the selected scenario—can all be managed.

On the other hand, when $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_K\}$ is only a subset of the feasible scenarios, then failure to simulate the optimal scenario can only be addressed in an

asymptotic sense, as described in the discussion on convergence in Sect. 8.3.1. The other two SO errors can, however, still be controlled. We will focus on the second one, selection error, since it is the most critical.

At the end of a SO search, the output data that have been generated are

$$Y_j(\mathbf{x}_i), j = 1, 2, \dots, n_i, i = 1, 2, \dots, K. \quad (8.29)$$

Since SO algorithms often revisit scenarios and accumulate observations on each visit, there is no reason to believe that the n_i 's will be equal, which limits the effectiveness of common random numbers.

Boesel et al. (2003) proposed the idea of “cleaning up” after an SO algorithm terminates by starting with the data (8.29) and doing as little additional simulation as possible to obtain a correct selection guarantee. The approach consists of two steps:

1. Apply a subset selection procedure to (8.29) to reduce (typically greatly) the number of scenarios that are competitive to be the best. The subset selection procedure described Sect. 8.4.2.1 is particularly useful for clean up because it can start from unequal sample sizes n_i , and it only requires that the sample means and sample variances from all K scenarios be retained, not all of the raw data.
2. Apply a selection of the best procedure to the scenarios in the subset. The procedure presented here could be used, but others described in Boesel et al. (2003) are easier to apply.

Boesel et al. (2003) showed that certain combinations of subset selection and selection of the best can be formally proven to deliver an overall correct selection guarantee. Xu et al. (2010) later showed that a precise estimation guarantee (the third SO error) can be added in some cases. But even applied informally clean up will greatly improve SO performance.

Remark 8.3. Instead of “cleaning up” one might hope to manage selection error at every iteration of the SO search, so that whenever the algorithm terminates there is a statistical guarantee that the current sample best scenario is the best of those simulated up to that iteration. Hong and Nelson (2007a) showed that this is possible, but computationally expensive unless K is small. Rapid progress in uncovering better and better scenarios in a large feasible space \mathbf{C} typically requires less simulation effort than is required to guarantee that the sample best scenario is the best at all iterations. This is why we recommend the clean up approach *after* the SO search.

8.4.3 Adaptive Random Search

We now consider simulation optimization problems such as (8.19) for which \mathbf{C} is discrete and finite, but contains far too many scenarios to simulate them all. Adaptive random search provides a framework for attacking such problems within which

asymptotic convergence can be proven, while still allowing for algorithms that aggressively pursue improving scenarios. We first describe adaptive random search at a high level, then provide a specific algorithm.

An adaptive random search algorithm is iterative, with the iterations indexed by $i = 1, 2, \dots$. On the i th iteration, the algorithm has a probability distribution $P_i(\mathbf{x})$ on the scenarios $\mathbf{x} \in C$. This distribution is used to sample one or more scenarios from C . The distribution $P_i(\cdot)$ may (and usually will) depend on its “memory” of some or all of the scenarios that have been simulated in previous iterations and perhaps the output data obtained from simulating them.

The distinction between *sampling* scenarios and *simulating* them is important: “Sampling” means randomly choosing a scenario \mathbf{x} from C , while “simulating” a scenario means generating output performance data. Adaptive random search algorithms have a simulation allocation rule that specifies how much simulation effort (e.g., how many replications) to expend on a sampled scenario, and also a value $V(\mathbf{x})$ for each simulated scenario. Most often $V(\mathbf{x}) = \hat{\theta}(\mathbf{x})$, its estimated performance, but it can be other measures such as the number of times \mathbf{x} has been visited by the search.

Initialize: Set the iteration counter to $i = 1$.

Sample: Choose an estimation set, which is a collection scenarios $E_i \subset C$ where some or all of the scenarios were chosen by sampling according to $P_i(\cdot)$ and others were retained from previous iterations.

Simulate: Apply the simulation allocation rule to simulate the scenarios $\mathbf{x} \in E_i$.

Evaluate: Update the value $V(\mathbf{x})$ for all $\mathbf{x} \in E_i$ and choose as $\hat{\mathbf{x}}_{i+1}^*$ the scenario with the best value $V(\mathbf{x})$.

Iterate: Update the algorithm memory, let $i = i + 1$ and go to **Sample**.

Notice that the algorithm contains no stopping rule, which makes sense for proving asymptotic convergence. When a specific algorithm is applied, stopping may occur when the simulation budget is exhausted or when progress appears to slow.

The sampling distribution $P_i(\cdot)$ and simulation allocation rule work together to guarantee asymptotic convergence. When no strong structural properties of $\theta(\mathbf{x})$ are known, then global convergence usually requires that all scenarios in C are simulated infinitely often in the limit. For local convergence a scenario and all of its neighbors must be simulated infinitely often for any scenario that appears to be locally optimal. See Andradóttir (1999, 2006b) and Hong and Nelson (2007a) for specific conditions that insure global and local convergence, respectively, of adaptive random search algorithms.

Three types of sampling distributions $P_i(\cdot)$ are common:

- A distribution that puts positive probability on a small number of feasible scenarios in a neighborhood of $\hat{\mathbf{x}}_i^*$. Typically, $P_i(\cdot)$ and the neighborhood structure connect C so that any scenario is reachable from any other scenario after a sufficient number of iterations.
- A distribution that puts positive probability on a “promising” subset of C that may be large or small, but is not necessarily a neighborhood of $\hat{\mathbf{x}}_i^*$. Typically

such distributions use memory in an intelligent way to concentrate the search. The algorithm described below is of this type.

- A distribution that puts positive probability on all of C . Typically the distribution changes as a function of the iteration and the memory, focusing probabilistically on promising regions of C .

To illustrate some of the key ideas in adaptive random search we describe the adaptive hyperbox algorithm (AHA) of Xu et al. (2012). This algorithm is applicable when the components of \mathbf{x} are integer-ordered decision variables and the neighborhood of \mathbf{x} is defined to be all feasible scenarios that differ from \mathbf{x} by ± 1 in one component. For instance, in two dimensions the scenario $\mathbf{x} = (3, 7)$ has neighbors $\{(2, 7), (4, 7), (3, 6), (3, 8)\}$. Under mild conditions AHA can be shown to converge to a locally optimal scenario (i.e., a scenario that is better than all of its feasible neighbors).

Here is the AHA implementation of adaptive random search:

Value: AHA uses the estimated objective function value $\hat{\theta}(\mathbf{x})$ accumulated over all output data generated for scenario \mathbf{x} . The key technical assumption is that $\hat{\theta}(\mathbf{x})$ converges with probability 1 to $\theta(\mathbf{x})$ as the amount of simulation effort (replications or run length) goes to infinity.

Memory: AHA remembers all scenarios that have been simulated during the course of the search and their estimated objective function values accumulated over all output data generated.

Distribution $P_i(\cdot)$: AHA puts positive probability on scenarios that are feasible and are in or on a hyperbox that surrounds $\hat{\mathbf{x}}_i^*$, the simulated scenario with the best estimated objective function value on iteration i . The hyperbox is defined by scenarios that have been simulated and are closest to $\hat{\mathbf{x}}_i^*$ in one or more coordinate. All other feasible scenarios have probability 0. A fixed number of scenarios are sampled on each iteration, and revisiting scenarios is allowed.

Estimation set E_i : For the estimation set AHA includes the current sample best $\hat{\mathbf{x}}_i^*$ and any scenarios sampled from $P_i(\cdot)$.

Simulation allocation rule: AHA requires that whenever a scenario is in the estimation set it must receive additional simulation (replications or run length).

Figure 8.4 illustrates how AHA works in a feasible region similar to the SAN optimization problem except that it is only in two dimensions. On the first iteration the three scenarios indicated by \bullet are randomly sampled from within the feasible C , all are simulated according to the simulation allocation rule, and the scenario indicated by $\hat{\mathbf{x}}_1^*$ is the sample best. Therefore, the other two \bullet scenarios are used to build the hyperbox. Scenarios for the next iteration are sampled from the intersection of the hyperbox with C ; this provides the \circ scenarios which are simulated along with $\hat{\mathbf{x}}_1^*$ according to the simulation allocation rule. Of these four scenarios, the one indicated by $\hat{\mathbf{x}}_2^*$ is now the sample best, so the scenarios that are closest in each coordinate direction are used to define the second hyperbox.

It is easy to see that if there was no simulation error ($\hat{\theta}(\mathbf{x}) = \theta(\mathbf{x})$), then the hyperbox would continue to shrink until it surrounds a single scenario with $\theta(\mathbf{x}^*)$ value less than or equal to all of its neighbors; that is, a locally optimal scenario.

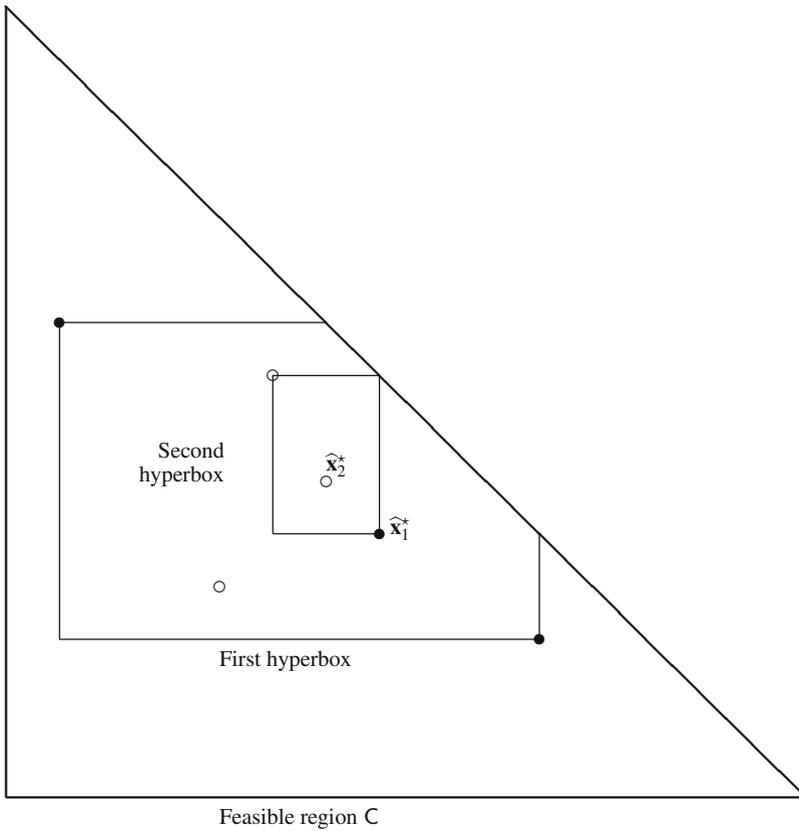


Fig. 8.4 First iteration (• scenarios) and second iteration (○ scenarios) of the adaptive hyperbox algorithm

Even with simulation error this will happen in the limit because the estimates $\hat{\theta}(\mathbf{x})$ of an apparently locally optimal scenario and its neighbors converge to their true expected values as the number of iterations (and therefore simulation effort) goes to infinity.

8.4.4 Searching in Improving Directions

We now consider simulation optimization problems such as (8.19) for which the objective function $\theta(\mathbf{x})$ is continuous and differentiable in \mathbf{x} , and C is (typically) a convex subset of \mathcal{R}^d . Many algorithms for deterministic nonlinear optimization problems search C for (at least locally) optimal scenarios by moving in improving directions based on the gradient of $\theta(\mathbf{x})$:

$$\nabla\theta(\mathbf{x}) = \left(\frac{\partial\theta(\mathbf{x})}{\partial x_1}, \frac{\partial\theta(\mathbf{x})}{\partial x_2}, \dots, \frac{\partial\theta(\mathbf{x})}{\partial x_d} \right).$$

For simulation optimization we can try to estimate $\nabla\theta(\mathbf{x})$, just as we estimate $\theta(\mathbf{x})$ itself, and then incorporate it into some sort of steepest descent algorithm such as stochastic approximation (e.g., Fu 2006). Gradient estimators are also useful in their own right for sensitivity analysis. Our approach in this section is to provide enough understanding of gradient estimation so that the reader can navigate the vast literature on this subject. More so than any other topic in this book, the mathematical conditions required for provable correctness of gradient estimators are difficult to explain and verify; good general references are Fu (2006) and L'Ecuyer (1990).

Recall that the basic output from the SAN optimization (8.20) can be expressed in two ways:

$$Y(\mathbf{x}) = \max\{A_1(x_1) + A_4(x_4), \\ A_1(x_1) + A_3(x_3) + A_5(x_5), \\ A_2(x_2) + A_5(x_5)\} \quad (8.30)$$

$$= \max\{-\ln(1 - U_1)x_1 - \ln(1 - U_4)x_4, \\ -\ln(1 - U_1)x_1 - \ln(1 - U_3)x_3 - \ln(1 - U_5)x_5, \\ -\ln(1 - U_2)x_2 - \ln(1 - U_5)x_5\} \quad (8.31)$$

where $(U_1, U_2, U_3, U_4, U_5)$ are the pseudorandom numbers needed to generate the exponentially distributed activity times (A_1, A_2, \dots, A_5) , and $\theta(\mathbf{x}) = E(Y(\mathbf{x}))$. We will illustrate key ideas for gradient estimation by estimating the first component of the gradient, $\partial\theta(\mathbf{x})/\partial x_1$, which is the partial derivative of the mean time to complete the project with respect to the mean of the first activity time, x_1 . Precisely the same ideas apply to the other components x_2, x_3, x_4, x_5 .

8.4.4.1 Finite Differences

Representation (8.31) and the definition of derivative suggests a natural approximation. Let

$$Y(\mathbf{x} + \Delta x_1) = \max\{-\ln(1 - U_1)(x_1 + \Delta x_1) - \ln(1 - U_4)x_4, \\ -\ln(1 - U_1)(x_1 + \Delta x_1) - \ln(1 - U_3)x_3 - \ln(1 - U_5)x_5, \\ -\ln(1 - U_2)x_2 - \ln(1 - U_5)x_5\}$$

which is the time to complete the project when the mean time for activity 1 is $x_1 + \Delta x_1$ (the term Δx_1 is used to denote the vector $(\Delta x_1, 0, 0, \dots, 0)^\top$). Then form the finite-difference (FD) estimator

$$\text{FD}(x_1) = \frac{Y(\mathbf{x} + \Delta x_1) - Y(\mathbf{x})}{\Delta x_1}. \quad (8.32)$$

Notice that we intentionally used common random numbers—the same $(U_1, U_2, U_3, U_4, U_5)$ for each scenario—which makes sense because $\text{FD}(x_1)$ estimates the difference between the mean time to complete the project for scenarios $\mathbf{x} + \Delta x_1$ and \mathbf{x} . Typically we would average $\text{FD}(x_1)$ estimators from across $n > 1$ replications to estimate $\partial\theta(\mathbf{x})/\partial x_1$.

FD gradient estimators are easy to understand and implement. However, there are substantial disadvantages: From a computational perspective, at least $n(d+1)$ simulations are required to estimate the gradient in all d coordinate directions; this can be a substantial burden if the optimization algorithm needs to repeatedly estimate the gradient as it searches. Also FD is biased because

$$E(\text{FD}(x_1)) = \frac{\theta(\mathbf{x} + \Delta x_1) - \theta(\mathbf{x})}{\Delta x_1}$$

while the derivative is the limit as $\Delta x_1 \rightarrow 0$. However, we cannot make Δx_1 too small because this causes numerical error as well as high variance since

$$\text{Var}[(Y(\mathbf{x} + \Delta x_1) - Y(\mathbf{x})) / \Delta x_1] = \text{Var}[Y(\mathbf{x} + \Delta x_1) - Y(\mathbf{x})] / \Delta x_1^2.$$

Common random numbers can provide substantial help with the variance problem and should virtually always be used (see Glasserman and Yao 1992).

8.4.4.2 Infinitesimal Perturbation Analysis

Suppose that we fix the pseudorandom numbers at $(u_1, u_2, u_3, u_4, u_5)$. If either $-\ln(1-u_1)x_1 - \ln(1-u_4)x_4$ or $-\ln(1-u_1)x_1 - \ln(1-u_3)x_3 - \ln(1-u_5)x_5$ is the longest path, then it will also be on the longest path at $\mathbf{x} + \Delta x_1$ for Δx_1 small enough. And since nothing is random, the limit $\Delta x_1 \rightarrow 0$ of the FD estimator is easy to calculate:

$$\begin{aligned} \lim_{\Delta x_1 \rightarrow 0} \frac{Y(\mathbf{x} + \Delta x_1) - Y(\mathbf{x})}{\Delta x_1} &= \lim_{\Delta x_1 \rightarrow 0} \frac{-\ln(1-u_1)(x_1 + \Delta x_1) - (-\ln(1-u_1)x_1)}{\Delta x_1} \\ &= -\ln(1-u_1) \\ &= \frac{-\ln(1-u_1)x_1}{x_1} = \frac{A_1(x_1)}{x_1}. \end{aligned} \quad (8.33)$$

If neither of these is the longest path, then the difference does not depend on Δx_1 and

$$\lim_{\Delta x_1 \rightarrow 0} \frac{Y(\mathbf{x} + \Delta x_1) - Y(\mathbf{x})}{\Delta x_1} = 0. \quad (8.34)$$

These two cases together comprise the infinitesimal perturbation analysis (IPA) gradient estimator

$$\text{IPA}(x_1) = \frac{A_1(x_1)}{x_1} I\{A_1(x_1) \text{ is on the longest path}\}. \quad (8.35)$$

Notice that $\text{IPA}(x_1)$ requires no additional simulation, just an additional calculation.

Another heuristic argument that ends up at the same place is this: To estimate $\partial\theta(\mathbf{x})/\partial x_1$ it makes some sense to use $\partial Y(\mathbf{x})/\partial x_1$. If A_1 is on the longest path, then we can apply the chain rule to get

$$\begin{aligned} \frac{\partial Y(\mathbf{x})}{\partial x_1} &= \frac{\partial Y(\mathbf{x})}{\partial A_1} \times \frac{\partial A_1(x_1)}{\partial x_1} \\ &= 1 \times -\ln(1 - U_1) = \frac{A_1(x_1)}{x_1}. \end{aligned}$$

If A_1 is not on the longest path then the derivative is again 0.

We have not established that $\text{IPA}(x_1)$ is a good gradient estimator (it is), but the reason it “works” is that, with probability 1, $Y(\mathbf{x})$ is continuous, differentiable and bounded in a neighborhood of \mathbf{x} for any fixed \mathbf{x} . Analogous IPA gradient estimators exist for x_2, x_3, x_4, x_5 , and they can all be computed simultaneously from a single replication. Of course, to estimate the gradient we average $\text{IPA}(x_1)$ over $n > 1$ replications.

More generally, the validity of IPA gradient estimation hinges on the validity of an interchange of differentiation and expectation:

$$\frac{\partial\theta(\mathbf{x})}{\partial x_i} = \frac{\partial E[Y(\mathbf{x})]}{\partial x_i} \stackrel{?}{=} E \left[\frac{\partial Y(\mathbf{x})}{\partial x_i} \right].$$

Technical conditions that allow this interchange are easy to state, but often not as easy to verify for a simulation model.

8.4.4.3 Likelihood Ratio

One problem with FD was the need for $d + 1$ simulations to estimate the full gradient. The following simple observation can get it down to 1.

Let $f_j(a_j|x_j) = \exp\{-a_j/x_j\}/x_j$ be the exponential distribution of the activity time A_j , $j = 1, 2, \dots, 5$ for a fixed scenario \mathbf{x} . Thus, the joint distribution of the activity times is given by $\prod_{j=1}^5 \exp\{-a_j/x_j\}/x_j$. Also let $g(\mathbf{a}) = \max\{a_1 + a_4, a_1 + a_3 + a_5, a_2 + a_5\}$ where $\mathbf{a} = (a_1, a_2, a_3, a_4, a_5)$. Consider the expectation of a weighted version of the project completion time $Y(\mathbf{x})$:

$$\begin{aligned} E \left[Y(\mathbf{x}) \frac{f_1(A_1|x_1 + \Delta x_1)}{f_1(A_1|x_1)} \right] &= E \left[Y(\mathbf{x}) \frac{\exp\{-A_1/(x_1 + \Delta x_1)\}/(x_1 + \Delta x_1)}{\exp\{-A_1/x_1\}/x_1} \right] \\ &= \int_0^\infty \cdots \int_0^\infty g(\mathbf{a}) \frac{f_1(a_1|x_1 + \Delta x_1)}{f_1(a_1|x_1)} \prod_{j=1}^5 f_j(a_j|x_j) \, d\mathbf{a} \end{aligned}$$

$$\begin{aligned}
&= \int_0^\infty \cdots \int_0^\infty g(\mathbf{a}) f_1(a_1|x_1 + \Delta x_1) \prod_{j=2}^5 f_j(a_j|x_j) \, d\mathbf{a} \\
&= E[Y(\mathbf{x} + \Delta x_1)] = \theta(\mathbf{x} + \Delta x_1).
\end{aligned}$$

Notice that $A_1 \sim$ exponential with mean x_1 , but by appropriately reweighting the output $Y(\mathbf{x})$ we can estimate $\theta(\mathbf{x} + \Delta x_1)$ from the observations generated under setting \mathbf{x} . The weight $f_1(A_1|x_1 + \Delta x_1)/f_1(A_1|x_1)$ is sometimes called the likelihood ratio because it can be interpreted as the relative likelihood of the observation A_1 if the mean was $x_1 + \Delta x_1$ as compared to x_1 . Therefore, from a single simulation replication we can get a finite-difference estimator:

$$\frac{Y(\mathbf{x}) \frac{f_1(A_1|x_1 + \Delta x_1)}{f_1(A_1|x_1)} - Y(\mathbf{x})}{\Delta x_1} = \frac{Y(\mathbf{x})}{f_1(A_1|x_1)} \times \frac{f_1(A_1|x_1 + \Delta x_1) - f_1(A_1|x_1)}{\Delta x_1}.$$

As we did with FD to motivate IPA, we can now take the $\lim_{\Delta x_1 \rightarrow 0}$. Notice that

$$\lim_{\Delta x_1 \rightarrow 0} \frac{f_1(A_1|x_1 + \Delta x_1) - f_1(A_1|x_1)}{\Delta x_1} = \frac{\partial f_1(A_1|x_1)}{\partial x_1}$$

by definition. Therefore, the likelihood ratio (LR) gradient estimator is

$$\text{LR}(x_1) = \frac{Y(\mathbf{x})}{f_1(A_1|x_1)} \frac{\partial f_1(A_1|x_1)}{\partial x_1} = Y(\mathbf{x}) \frac{\partial \ln f_1(A_1|x_1)}{\partial x_1}. \quad (8.36)$$

For the exponential distribution of activity 1

$$\frac{\partial \ln f_1(A_1|x_1)}{\partial x_1} = \frac{\partial}{\partial x_1} \left(-\frac{A_1}{x_1} - \ln x_1 \right) = \frac{A_1}{x_1^2} - \frac{1}{x_1} = \frac{A_1 - x_1}{x_1^2}.$$

To estimate the gradient we average $\text{LR}(x_1)$ over $n > 1$ replications. Clearly we can do the same thing for x_2, x_3, x_4, x_5 . Again there are technical conditions that need to be satisfied for this to be a valid gradient estimator, and they are satisfied in this example. For LR to be applicable x must be the parameter of an input distribution.

8.4.4.4 Weak Derivative Gradient Estimators

We motivate the final gradient estimator by wishful thinking. Recall that $\text{FD}(x_1)$ comes from simulating A_1 from distribution f_1 with mean $x_1 + \Delta x_1$ and also with mean x_1 , and then computing $(Y(\mathbf{x} + \Delta x_1) - Y(\mathbf{x}))/\Delta x_1$. $\text{FD}(x_1)$ has expectation $(\theta(\mathbf{x} + \Delta x_1) - \theta(\mathbf{x}))/\Delta x_1$, which is not quite what we want. The wishful thinking is that there exist two other distributions for A_1 , say $f_1^{(1)}$ and $f_1^{(2)}$, and a scaling constant $\Delta(x_1)$, that are just right, meaning

$$\mathbb{E} \left[\frac{g(A_1^{(1)}, A_2, A_3, A_4, A_5) - g(A_1^{(2)}, A_2, A_3, A_4, A_5)}{\Delta(x_1)} \right] = \frac{\partial \theta(\mathbf{x})}{\partial x_1}.$$

That is, we do two simulations, like FD, generating A_2, A_3, A_4, A_5 from their respective exponential distributions with means x_2, x_3, x_4, x_5 , but we generate A_1 from $f_1^{(1)}$ for the first simulation, from $f_1^{(2)}$ for the second one, and then take a scaled finite difference.

Fortunately, it turns out that a magic triple $(\Delta(x), f^{(1)}, f^{(2)})$ exists for many common distributions (see Fu 2006, Table 1, noting that his $c = 1/\Delta(x)$). These are called weak derivatives (WD). For instance, for the exponential distribution with mean x a weak derivative is

$$\left(\Delta(x), f^{(1)}(a), f^{(2)}(a) \right) = \left(x, \frac{a}{x^2} e^{-a/x}, \frac{1}{x} e^{-a/x} \right).$$

Note that $ae^{-a/x}/x^2$ is an Erlang-2 distribution with mean $2x$.

Think of WD as a FD gradient estimator with no bias. As with FD, common random numbers should be used for the two (or in general $2d$) simulations.

8.4.4.5 Choice of Gradient Estimator

The appropriate gradient estimator depends very much on specifics of the simulation. The following discussion provides some guidance, but more detailed references like Fu (2006) and L'Ecuyer (1990) should be consulted.

FD is always applicable provided $\theta(\mathbf{x})$ is differentiable and $\hat{\theta}(\mathbf{x})$ has finite variance in a neighborhood of \mathbf{x} . Of course, FD requires a minimum of $d + 1$ simulations and a choice of Δx . Common random numbers should always be used. At the cost of $2d$ simulations a better estimator uses central differences

$$\text{CD}(x_1) = \frac{Y(\mathbf{x} + \Delta x_1) - Y(\mathbf{x} - \Delta x_1)}{2\Delta x_1}.$$

If d is small or gradient estimates are only needed for a few scenarios \mathbf{x} then FD and CD are attractive.

IPA and LR have the advantage of saving computational effort since they can be computed from a single simulation no matter what the dimension of \mathbf{x} . IPA tends to have smaller variance, but the conditions under which it is appropriate are sometimes difficult to verify and IPA can require more significant programming to implement (notice that in the SAN optimization IPA necessitated determining which activities were on the longest path). IPA gradient estimators are most often available when the simulation output itself can be represented—either directly or via the chain rule—as a differentiable function of x_1 .

LR gradient estimators, on the other hand, are most often available when x_1 can be represented as a parameter of an input distribution. When LR gradient estimators

exist, then computing them only requires reweighting the output of interest with a weight that is a function of the input distribution. LR gradient estimators tend to have larger variance than IPA, especially when x is the parameter of more than one random variate per replication because the weight is a product of the input distribution for each instance (e.g., x is the mean service time in a queue that simulates 1,000 customers per replication so the weight will involve the product of the service time distribution 1,000 times).

Like LR, WD gradient estimators can be considered when x is a parameter of an input distribution. WD gradient estimators do not save simulations relative to FD, and will be most useful when the input whose distribution has parameter x appears only once per replication, as it did in the SAN simulation. WD gradient estimators tend to have low variance when common random numbers are employed.

Just like the simulation response estimator $\hat{\theta}(\mathbf{x}; T, n, \mathbf{U})$ itself, the variance of the gradient estimator depends on the number of replications and, for a steady-state simulation, the run length. The number of replications or run length required to get a precise estimate of $\theta(\mathbf{x})$ may or may not be sufficient to obtain a precise estimate of its gradient.

8.4.4.6 Steepest Descent

Gradient-based search for SO is beyond the scope of this book, other than to describe the simplest method known as *stochastic approximation*, which is a method of steepest descent: Starting from an initial scenario \mathbf{x}_0 , implement the recursion

$$\mathbf{x}_{i+1} = \mathbf{x}_i - a_i \widehat{\nabla} \theta(\mathbf{x}_i) \quad (8.37)$$

where a_i is a sequence such that $a_i \rightarrow 0$, but $\sum_{i=1}^{\infty} a_i = \infty$; setting $a_i = a/i$ for a positive constant a is a common choice. The intuition is that when $\nabla \theta(\mathbf{x}_i) = 0$ a stationary point has been reached that is at least a local minimum, while having $a_i \rightarrow 0$ (but not too fast) mutes the impact of variability in the gradient estimator $\widehat{\nabla} \theta(\mathbf{x}_i)$.

Stochastic approximation can be sensitive to the choice of sequence a_i , so it is often made adaptive. A second issue is that on some iterations \mathbf{x}_{i+1} may end up outside the feasible region C ; then there must be some way to project it back into C .

For instance, the constraints for the SAN optimization are

$$\begin{aligned} \sum_{j=1}^5 c_j(\tau_j - x_j) &\leq b \\ x_j &\geq \ell_j, \quad j = 1, 2, 3, 4, 5. \end{aligned}$$

If $\sum_{j=1}^5 c_j(\tau_j - x_{i+1,j}) > b$, then we might instead use \mathbf{x}'_{i+1} which is the perpendicular projection of \mathbf{x}_{i+1} onto the plane $\sum_{j=1}^5 c_j(\tau_j - x_j) = b$, provided it also satisfies the lower bound constraints. Exercise 19 asks you to work this out.

8.4.5 Optimizing the Sample Average Problem

In SO our goal is to minimize $\theta(\mathbf{x})$ using a simulation-based estimator $\widehat{\theta}(\mathbf{x}; T, n, \mathbf{U})$ of the objective function; we will assume that the replication stopping time T is not a factor and drop it from the notation. Given fixed values for the number of replications n and, most importantly, the random numbers $\mathbf{U} = \mathbf{u}$, consider solving instead the optimization problem

$$\begin{aligned} \min \widehat{\theta}(\mathbf{x}; n, \mathbf{U} = \mathbf{u}) \\ \mathbf{x} \in C. \end{aligned} \quad (8.38)$$

In words, the objective of the optimization problem in (8.38) is to minimize the value of the estimator $\widehat{\theta}(\mathbf{x}; n, \mathbf{U} = \mathbf{u})$ as a function of \mathbf{x} . With the random numbers fixed, this is a deterministic optimization problem and it seems plausible that if n is large enough then the optimal solution to (8.38) should be close to the optimal solution of the SO problem. This approach is called *sample average approximation (SAA)*.

For instance, the output from replication j of the SAN simulation can be expressed as

$$\begin{aligned} y_j(\mathbf{x}) &= \max\{a_{1j}(x_1) + a_{4j}(x_4), \\ &\quad a_{1j}(x_1) + a_{3j}(x_3) + a_{5j}(x_5), \\ &\quad a_{2j}(x_2) + a_{5j}(x_5)\} \\ &= \max\{-\ln(1 - u_{1j})x_1 - \ln(1 - u_{4j})x_4, \\ &\quad -\ln(1 - u_{1j})x_1 - \ln(1 - u_{3j})x_3 - \ln(1 - u_{5j})x_5, \\ &\quad -\ln(1 - u_{2j})x_2 - \ln(1 - u_{5j})x_5\} \end{aligned}$$

where we have made Y, A and U lowercase to indicate that there are no random variables left after the random numbers are fixed. The objective of the corresponding SAA problem is

$$\min \widehat{\theta}(\mathbf{x}; n, \mathbf{U} = \mathbf{u}) = \min \frac{1}{n} \sum_{j=1}^n y_j(\mathbf{x}).$$

In many cases the SAA problem is a difficult nonlinear optimization problem without apparent structure and therefore cannot be readily solved. But in some SO problems, such as the SAN optimization (8.20), solving the SAA problem is easy. In fact, if we replace the max operator by three inequality constraints then the SAA problem can be formulated as a linear program:

$$\begin{aligned} \min \quad & \frac{1}{n} \sum_{j=1}^n y_j \\ & y_j \geq -\ln(1 - u_{1j})x_1 - \ln(1 - u_{4j})x_4 \\ & y_j \geq -\ln(1 - u_{1j})x_1 - \ln(1 - u_{3j})x_3 - \ln(1 - u_{5j})x_5 \end{aligned} \quad (8.39)$$

$$\begin{aligned}
 y_j &\geq -\ln(1 - u_{2j})x_2 - \ln(1 - u_{5j})x_5, \quad j = 1, 2, \dots, n \\
 b &\geq \sum_{k=1}^5 c_k(\tau_k - x_k) \\
 x_k &\geq \ell_k, \quad k = 1, 2, \dots, 5.
 \end{aligned}$$

This linear program has $n + 5$ decision variables $y_1, y_2, \dots, y_n, x_1, x_2, \dots, x_5$ and $3n + 6$ constraints (remember that all of the $-\ln(1 - u_{ij})$ terms are fixed constants, and therefore are coefficients in the linear program). Thus, the SAA problem can be solved even for a very large number of replications, n .

Let $\widehat{\mathbf{x}}_n^*$ be the optimal solution to the SAA problem (8.38) based on n replications. Then what matters is whether $\theta(\widehat{\mathbf{x}}_n^*)$ is close to $\theta(\mathbf{x}^*)$, and in particular does $\theta(\widehat{\mathbf{x}}_n^*) \rightarrow \theta(\mathbf{x}^*)$ as $n \rightarrow \infty$? The various conditions that guarantee convergence are quite technical, so we present just one set here. A comprehensive reference is Shapiro et al. (2009).

Suppose that for any $\mathbf{x} \in \mathcal{C}$ we have $\widehat{\theta}(\mathbf{x}; n, \mathbf{U}) \rightarrow \theta(\mathbf{x})$ with probability 1 as the number of replications $n \rightarrow \infty$; that is, $\widehat{\theta}(\mathbf{x}; n, \mathbf{U})$ satisfies a strong law of large numbers as is often the case for averages across replications (see Sect. 5.2.2). From here on we will assume that $\widehat{\theta}(\mathbf{x}; n, \mathbf{U})$ is an average across i.i.d. replications and $E[\widehat{\theta}(\mathbf{x}; n, \mathbf{U})] = \theta(\mathbf{x})$.

Although pointwise convergence might seem like enough, it is insufficient for convergence of the optimal value $\theta(\widehat{\mathbf{x}}_n^*) \rightarrow \theta(\mathbf{x}^*)$ in general because $\widehat{\mathbf{x}}_n^*$ is not a fixed value of \mathbf{x} , but rather the solution to the SAA problem over *all* $\mathbf{x} \in \mathcal{C}$. A more comprehensive form of convergence is needed, in particular $\widehat{\theta}(\mathbf{x}; n, \mathbf{U}) \rightarrow \theta(\mathbf{x})$ *uniformly* for all $\mathbf{x} \in \mathcal{C}$.

Uniform convergence means that, for any $\varepsilon > 0$, with probability 1 there exists an n' such that

$$\sup_{\mathbf{x} \in \mathcal{C}} \left| \widehat{\theta}(\mathbf{x}; n; \mathbf{U}) - \theta(\mathbf{x}) \right| \leq \varepsilon$$

for all $n \geq n'$. For pointwise convergence this n' can depend on \mathbf{x} ; for uniform convergence it has to be independent of \mathbf{x} , although it can depend on the particular sequence of random numbers \mathbf{U} . What it means is that when the number of replications is large enough the estimator $\widehat{\theta}(\mathbf{x}; n; \mathbf{U})$ is uniformly close to $\theta(\mathbf{x})$ for all \mathbf{x} . As a practical matter, if the estimator $\widehat{\theta}(\mathbf{x}; n; \mathbf{U})$ is continuous in \mathbf{x} with probability 1, and $\theta(\mathbf{x})$ is bounded, then we will have uniform convergence. Verifying this has to be done problem by problem.

When it applies SAA can be a powerful SO method because it turns a stochastic problem into a deterministic one to which we can apply sophisticated linear and nonlinear optimization techniques. A drawback is that the computational requirements to solve the SAA problem can increase substantially in n , and we want n large to obtain a good approximation.

8.4.6 Stochastic Constraints

Recall again the formulation of the simulation optimization (SO) problem that opened this section:

$$\begin{aligned} \min \theta(\mathbf{x}) \\ \mathbf{x} \in C . \end{aligned}$$

What makes it a SO problem is the need to estimate $\theta(\mathbf{x})$ using a stochastic simulation. We have assumed that for any scenario \mathbf{x} we can determine with certainty whether \mathbf{x} is feasible ($\mathbf{x} \in C$) or infeasible ($\mathbf{x} \notin C$). For many practical problems this is not possible. For instance, in the service center simulation of Sect. 4.6 the objective was to minimize staffing costs subject to two constraints:

$$\begin{aligned} \text{Percentage of simple orders entered in } \leq 10 \text{ min is at least } 96\% \\ \text{Percentage of special orders entered in } \leq 10 \text{ min is at least } 80\% \end{aligned}$$

Given a staffing scenario \mathbf{x} , satisfaction of these constraints has to be estimated.

For the service center the objective function (minimize staffing cost) was deterministic, but more generally we have to estimate it as well. So we expand the SO formulation to include stochastic constraints, specifically

$$\min \theta(\mathbf{x}) \tag{8.40}$$

$$\mathbf{x} \in C$$

$$c_\ell(\mathbf{x}) \geq q_\ell, \ell = 1, 2, \dots, v \tag{8.41}$$

where $\theta(\mathbf{x}), c_1(\mathbf{x}), c_2(\mathbf{x}), \dots, c_v(\mathbf{x})$ must all be estimated. For the purposes of this section we assume that for any scenario \mathbf{x} we can observe simulation outputs $Y(\mathbf{x}), C^{(1)}(\mathbf{x}), C^{(2)}(\mathbf{x}), \dots, C^{(v)}(\mathbf{x})$ with the property that

$$\begin{aligned} E(Y(\mathbf{x})) &= \theta(\mathbf{x}) \\ E(C^{(\ell)}(\mathbf{x})) &= c_\ell(\mathbf{x}), \ell = 1, 2, \dots, v \end{aligned}$$

so that it makes sense to estimate $\theta(\mathbf{x})$ and $c_\ell(\mathbf{x})$ by sample averages.

For example, in the service center simulation let $C^{(1)}(\mathbf{x})$ be the observed fraction of simple orders arriving between 8 a.m. and 4 p.m. and that are entered in ≤ 10 min, and let $C^{(2)}(\mathbf{x})$ be the corresponding fraction for special orders, under staffing scenario \mathbf{x} . Then the stochastic constraints are

$$c_1(\mathbf{x}) = E(C^{(1)}(\mathbf{x})) \geq 0.96$$

$$c_2(\mathbf{x}) = E(C^{(2)}(\mathbf{x})) \geq 0.80.$$

To simplify the presentation we will focus on the case of $v = 1$ stochastic constraint $c(\mathbf{x}) \geq q$ that is estimated by the simulation output $C(\mathbf{x})$.

Constraints nearly always represent limits on our ability to minimize the objective $\theta(\mathbf{x})$; stated differently, if the constraints could be relaxed or removed then scenarios could be chosen to make $\theta(\mathbf{x})$ even smaller. In the service center we could certainly assign fewer staff if, say, only 50% of the orders needed to be entered in 10 min or less. Therefore, at the optimal scenario \mathbf{x}^* we expect the constraints to be tight ($c(\mathbf{x}^*) = q$) or nearly so. *A key message of this section is that using simulation-based estimates to decide if $c(\mathbf{x}) \geq q$ when the constraint is tight, or nearly so, is a more difficult problem than estimating the value of $\theta(\mathbf{x})$.*

We will not recommend specific algorithms for dealing with stochastic constraints; instead, we make clear the difficulties that arise when including them. This section is somewhat advanced and can be skipped without loss of continuity.

8.4.6.1 Feasibility

When there are a relatively small number of scenarios to consider, then a reasonable strategy is to first decide which of them are feasible, and next conduct a simulation experiment to find the best among the feasible ones. How hard is it to decide if a scenario \mathbf{x} is feasible?

Suppose that for scenario \mathbf{x} we run n i.i.d. replications to obtain observations of the constrained output $C_1(\mathbf{x}), C_2(\mathbf{x}), \dots, C_n(\mathbf{x})$. In the service center simulation this might be the observed fraction of simple orders entered in ≤ 10 min for staffing level \mathbf{x} from n replications (days). We then decide constraint satisfaction by checking

$$\bar{C}(\mathbf{x}) \stackrel{?}{\geq} q.$$

Let $\sigma^2(\mathbf{x}) = \text{Var}(C(\mathbf{x}))$ and assume the usual case that $0 < \sigma^2(\mathbf{x}) < \infty$. Then by the central limit theorem (Sect. 5.2.2)

$$\Pr\{\bar{C}(\mathbf{x}) \geq q\} = \Pr\left\{\frac{\sqrt{n}(\bar{C}(\mathbf{x}) - c(\mathbf{x}))}{\sigma(\mathbf{x})} \geq \frac{\sqrt{n}(q - c(\mathbf{x}))}{\sigma(\mathbf{x})}\right\} \xrightarrow{n \rightarrow \infty} \Pr\{N(0, 1) \geq \lambda\} \quad (8.42)$$

where $N(0, 1)$ represents a standard normal random variable. There are three cases for λ :

1. If \mathbf{x} is infeasible then $q - c(\mathbf{x}) > 0$, so λ is ∞ and (8.42) is 0. This is what we want.
2. If \mathbf{x} is feasible but not tight then $q - c(\mathbf{x}) < 0$, so λ is $-\infty$ and (8.42) is 1. This is also what we want.
3. If \mathbf{x} is tight then $q - c(\mathbf{x}) = 0$, so λ is 0 and (8.42) is $1/2$. *This means that no matter how much simulation we do, we cannot determine with certainty if a constraint is tight, which implies that it will also be difficult when the constraint is close to being tight.*

When checking feasibility of stochastic constraints there are two types of errors: declaring a scenario feasible when it is infeasible, and declaring a scenario infeasible when it is feasible. Often, however, the seriousness of the error depends on the difference $q - c(\mathbf{x})$. In the service center if we declare that a staffing policy \mathbf{x} is feasible when in fact $c_1(\mathbf{x}) = 0.77$ then this is certainly a more serious error than declaring \mathbf{x}' feasible when $c_1(\mathbf{x}') = 0.94$.

A practical way to address this difficulty is to expand the concept of feasibility. Choose two additional cut-offs, q^- and q^+ such that $q^- < q < q^+$. If the value of $c(\mathbf{x}) \geq q^+$ then scenario \mathbf{x} is *desirable*; if $c(\mathbf{x}) \leq q^-$ then scenario \mathbf{x} is *unacceptable*; while if $q^- < c(\mathbf{x}) < q^+$ then scenario \mathbf{x} is *acceptable*. We then run experiments that allow us to be confident that the scenarios we declare feasible include all of the desirable ones and some or all of the acceptable ones, but does not include unacceptable scenarios. By allowing the slack to declare acceptable scenarios to be feasible—whether or not they are actually feasible—we avoid the difficulty with tight constraints.

Andradóttir and Kim (2010) and Batur and Kim (2010) provide formal procedures to accomplish this. Informally, we can estimate $c(\mathbf{x})$ precisely enough, as measured by a confidence interval $\bar{C}(\mathbf{x}) \pm H$, so that we can feel confident that \mathbf{x} is in the desirable to acceptable range (we do not have to separate them), or in the unacceptable range; we retain the former and discard the latter.

Returning to the service center example, if the requirement of at least 96% of all simple orders being entered in ≤ 10 min is not a firm requirement, then we might set $q^- = 0.92$ and $q^+ = 0.98$, for instance. This means that we want to be highly confident of declaring staffing scenarios that achieve 98% entry as feasible, to discard scenarios that are worse than 92%, and are willing to consider scenarios that achieve greater than a 92% but necessarily a 96% service level.

8.4.6.2 Constraint-Guided Search

Here we consider SO problems that require a search, rather than exhausting all possible scenarios. One approach is to represent violation of the stochastic constraints as a penalty. Two cases in this setting are particularly important.

Suppose that it is physically possible to violate the constraints, but the greater the violation the less desirable the scenario is. This is certainly true in the service center example. We might then reformulate the SO problem as

$$\begin{aligned} \min \quad & \theta(\mathbf{x}) + \sum_{\ell=1}^v p_{\ell}(q_{\ell} - c_{\ell}(\mathbf{x})) \\ & \mathbf{x} \in C \end{aligned} \tag{8.43}$$

where $p_{\ell}(\cdot)$ is a penalty function that is 0 if $c_{\ell}(\mathbf{x}) \geq q_{\ell}$ and increasing in $q_{\ell} - c_{\ell}(\mathbf{x})$ otherwise. For instance,

$$p_{\ell}(q_{\ell} - c_{\ell}(\mathbf{x})) = \beta_{\ell} \max\{0, q_{\ell} - c_{\ell}(\mathbf{x})\}.$$

This converts the SO problem (8.43) with stochastic constraints into a SO problem with deterministic constraints to which the methods described elsewhere in this section can be applied. Clearly the choice of the penalty function is important; in the service center problem it should convert underachieving the service level into a dollar cost (since the primary objective is staffing cost). Beyond that, there is a subtle statistical issue that should be noted.

Returning to the case of $v = 1$ constraint, suppose we have n replications of the simulation at scenario \mathbf{x} . How do we estimate the Objective (8.43)? A natural (and valid) estimator is

$$\bar{Y}(\mathbf{x}) + p(q - \bar{C}(\mathbf{x})). \quad (8.44)$$

If the penalty function $p(\cdot)$ is continuous around $q - c(\mathbf{x})$ then this estimator converges to the desired objective $\theta(\mathbf{x}) + p(q - c(\mathbf{x}))$ as $n \rightarrow \infty$ by the strong law of large numbers and the continuous mapping theorem (see Sect. 5.2.4).

However, one might be tempted to compute the penalty on each replication and use

$$\frac{1}{n} \sum_{j=1}^n [Y_j(\mathbf{x}) + p(q - C_j(\mathbf{x}))]. \quad (8.45)$$

This estimator will converge to

$$E[Y(\mathbf{x}) + p(q - C(\mathbf{x}))] \neq \theta(\mathbf{x}) + p(q - c(\mathbf{x}))$$

unless $p(\cdot)$ is linear. And since the penalty function is nonnegative, even a strictly feasible scenario may have a positive penalty in the limit.

Now suppose that we really want to enforce feasibility, at least asymptotically. Then instead of penalizing the objective function as in (8.43) we penalize the estimated value of the scenario. To describe this approach, recall that search algorithms work iteratively attempting to find better and better scenarios as the number of iterations increase. Let $i = 0, 1, 2, \dots$ be the iteration index, and let $n_i(\mathbf{x})$ be the total number of replications obtained from scenario \mathbf{x} through iteration i (search algorithms can and often do revisit scenarios). Then in the case of a single constraint we might define the estimated value of scenario \mathbf{x} through iteration i to be

$$\hat{\theta}(\mathbf{x}) = \bar{Y}(\mathbf{x}) + \beta_i \max\{0, q - \bar{C}(\mathbf{x})\}$$

where the averages include all observations $j = 1, 2, \dots, n_i(\mathbf{x})$. Notice that the coefficient β_i depends on the iteration i (it might also depend on $n_i(\mathbf{x})$) because we want the penalty to grow as the search progresses so that, at least asymptotically, infeasible scenarios are not competitive. The point to be made here is that one needs to be careful about using the simple prescription that $\beta_i \rightarrow \infty$ as $i \rightarrow \infty$. Here is why:

Suppose that the number of replications at scenario \mathbf{x} , $n_i(\mathbf{x})$, goes to infinity as the number of iterations goes to infinity. What happens to the penalty $\beta_i \max\{0, q - \bar{C}(\mathbf{x})\}$? It makes intuitive sense, and can be proven formally using the strong law of large numbers (see Sect. 5.2.2), that if \mathbf{x} is infeasible then $\beta_i \max\{0, q - \bar{C}(\mathbf{x})\}$

$\xrightarrow{a.s.} \infty$, while if \mathbf{x} is feasible but not tight then $\beta_i \max\{0, q - \bar{C}(\mathbf{x})\} \xrightarrow{a.s.} 0$, even though $\beta_i \rightarrow \infty$.

But this is misleading. Consider the probability of a penalty larger than, say, $\delta > 0$:

$$\begin{aligned} & \Pr\{\beta_i \max\{0, q - \bar{C}(\mathbf{x})\} > \delta\} \\ & \geq \Pr\{\beta_i (q - \bar{C}(\mathbf{x})) > \delta\} \\ & = \Pr\left\{\bar{C}(\mathbf{x}) - q < -\frac{\delta}{\beta_i}\right\} \\ & = \Pr\left\{\frac{\sqrt{n_i(\mathbf{x})}(\bar{C}(\mathbf{x}) - c(\mathbf{x}))}{\sigma} < \frac{\sqrt{n_i(\mathbf{x})}}{\sigma}\left(q - c(\mathbf{x}) - \frac{\delta}{\beta_i}\right)\right\}. \end{aligned}$$

Using a central limit theorem argument similar to Sect. 8.4.6.1 the random variable on the left-hand side is asymptotically $N(0, 1)$. Consider three cases for the right-hand side:

1. If \mathbf{x} is infeasible then we want this probability to get large quickly. Since $q - c(\mathbf{x}) > 0$ and $\delta > 0$, this means we would like β_i to become large quickly (big penalty), so $\beta_i \rightarrow \infty$ is fine.
2. If \mathbf{x} is feasible but not tight, then we want this probability to get small quickly (no penalty). Since $q - c(\mathbf{x}) < 0$ and $\delta > 0$, we would like β_i to remain small and ideally go to 0.
3. If \mathbf{x} is tight then we also want this probability to get small quickly (no penalty). Since $q - c(\mathbf{x}) = 0$ we *need* β_i to go to 0 to avoid a penalty.

Thus, contrary to the simple prescription we actually want $\beta_i \rightarrow \infty$ only for infeasible scenarios (whose identities we do not know in advance), suggesting that penalties should adapt to the observed (in)feasibility.

8.5 Variance Reduction Using Control Variates

In this final section we consider the choice of estimator $\hat{\theta}(\mathbf{x}; T, n, \mathbf{U})$ as part of experiment design. While sample means are often appropriate, sometimes we can do better.

We have emphasized measuring and controlling estimation error through the number of replications or runlength, and for many simulations that approach is sufficient. But estimation error typically declines as $1/\sqrt{\text{effort}}$, where “effort” is measured in replications or run length or both. Thus, the decline is slow and there are problems for which the simulation experiment is too computationally expensive to attain the level of error we want with the amount of simulation effort we can afford or have time to spend. Examples include:

- Estimating the performance of thousands of alternatives so that it is not possible to spend substantial simulation effort on each of them.
- Estimating performance measures associated with rare events so that very long simulations are required to observe even one occurrence of the event.
- Estimating performance parameters for which very little error can be tolerated when the outputs themselves are highly variable.

Variance reduction techniques (VRTs) are strategies for reducing estimation error without a corresponding increase in simulation effort. Common random numbers, described in Sect. 8.3, is a VRT since it reduces the variance of the estimated difference between two systems' performance measures relative to using independent simulations. "Variance reduction" is always relative to what the variance would have been had we not used the VRT.

VRTs often have to be carefully tailored to the specific simulation of interest. Good references are Bratley et al. (1987); Glasserman (2004); Asmussen and Glynn (2007) and Law (2007). Here we describe a VRT called *control variates* that, like common random numbers, is widely applicable.

Recall that common random numbers applies when the goal is to estimate $\mu(x_1) - \mu(x_2) = E[Y(x_1) - Y(x_2)]$, the difference in the expected performance of alternatives x_1 and x_2 ; it is based on the knowledge that

$$\text{Var}[Y(x_1) - Y(x_2)] = \text{Var}(Y(x_1)) + \text{Var}(Y(x_2)) - 2\text{Cov}(Y(x_1), Y(x_2))$$

so that forcing $\text{Cov}(Y(x_1), Y(x_2)) > 0$ by using common random numbers reduces variance.

Control variates apply when the goal is to estimate $\mu_Y = E(Y)$ for some simulation output Y ; it can be motivated by a different fact about variance: Let C be some other random variable that is also observable in the simulation. Then

$$\mu_Y = E(Y) = E[E(Y|C)] \tag{8.46}$$

$$\sigma_Y^2 = \text{Var}(Y) = \text{Var}[E(Y|C)] + E[\text{Var}(Y|C)]. \tag{8.47}$$

Recall that $E(Y|C)$, the conditional expected value of output Y given C , is itself a random variable. In the double-expectation result (8.46) and the variance-decomposition result (8.47), the inner expectation is with respect to the conditional distribution of Y given C , while the outer is with respect to the distribution of C . Together these results imply that the random variable $E(Y|C)$ is an unbiased estimator of μ_Y , it has no larger variance than Y , and it likely has smaller variance since we expect $E[\text{Var}(Y|C)] > 0$. Intuitively, the stronger the association between Y and C , the greater the variance reduction. Using $E(Y|C)$ instead of Y to estimate μ_Y is a VRT. However, in most real problems we will not know a conditional expected value $E(Y|C)$ that we can exploit, so we use these results in a different way.

The equalities (8.46) and (8.47) imply that we can represent the simulation output Y as follows (where we have listed the variance contribution below each term)

$$\underbrace{Y}_{\sigma_Y^2} = \underbrace{E(Y|C)}_{\text{Var}[E(Y|C)]} + \underbrace{\varepsilon}_{E[\text{Var}(Y|C)]}$$

where $\varepsilon = Y - E(Y|C)$, so that it has expected value 0, variance $E[\text{Var}(Y|C)]$ and ε is independent of $E(Y|C)$ (independence is not obvious but can be shown). Because we observe Y and C we have some hope of estimating the relationship $E(Y|C)$ and therefore (mostly) removing it as a contributor to variance. Control variates does this by using a linear approximation

$$E(Y|C) = \beta_0 + \beta_1(C - \mu_C) \quad (8.48)$$

where $\mu_C = E(C)$ and is *known*, and β_0 and β_1 are constants. Then it follows immediately that $\beta_0 = \mu_Y$, since

$$\mu_Y = E(Y) = E[E(Y|C)] = E[\beta_0 + \beta_1(C - \mu_C)] = \beta_0.$$

This gives the control-variate model

$$Y = \mu_Y + \beta_1(C - \mu_C) + \varepsilon. \quad (8.49)$$

Is the representation (8.49) even plausible? Here is one justification: If the joint distribution of (Y, C) is bivariate normal,

$$\begin{pmatrix} Y \\ C \end{pmatrix} \sim \text{BVN} \left(\begin{bmatrix} \mu_Y \\ \mu_C \end{bmatrix}, \begin{bmatrix} \sigma_Y^2 & \rho\sigma_Y\sigma_C \\ \rho\sigma_Y\sigma_C & \sigma_C^2 \end{bmatrix} \right) \quad (8.50)$$

where ρ is the correlation between Y and C , then it is known that

$$E(Y|C) = \mu_Y + \beta_1^*(C - \mu_C)$$

where

$$\beta_1^* = \frac{\text{Cov}(Y, C)}{\text{Var}(C)} = \frac{\rho\sigma_Y\sigma_C}{\sigma_C^2}$$

(see, e.g., Kotz et al. 2000). If we remove all of the variance due to $E(Y|C)$, then the leftover variance is $E[\text{Var}(Y|C)] = (1 - \rho^2)\sigma_Y^2 \leq \sigma_Y^2 = \text{Var}(Y)$, with the variance decreasing the larger the value of the squared correlation between Y and C .

Of course, we cannot count on our outputs being jointly normal. But in a simulation experiment we observe $(Y_i, C_i), i = 1, 2, \dots, n$, across n replications and base our analysis on (\bar{Y}, \bar{C}) . Then a multivariate version of the central limit theorem implies that under mild conditions the sample means (\bar{Y}, \bar{C}) will tend to be approximately bivariate normal if n is large enough. This provides one justification for the approximation (8.49).

When we have replications, the form of (8.49) suggests estimating $\beta_0 = \mu_Y$ via least-squares regression. The regression formulation is

$$\mathbf{Y} = \begin{pmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_n \end{pmatrix} = \begin{pmatrix} 1 & C_1 - \mu_C \\ 1 & C_2 - \mu_C \\ \vdots & \vdots \\ 1 & C_n - \mu_C \end{pmatrix} \begin{pmatrix} \beta_0 \\ \beta_1 \end{pmatrix} + \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_n \end{pmatrix} = \mathbf{C}\boldsymbol{\beta} + \boldsymbol{\varepsilon}.$$

This gives the least-squares estimators

$$\begin{pmatrix} \widehat{\beta}_0 \\ \widehat{\beta}_1 \end{pmatrix} = \widehat{\boldsymbol{\beta}} = (\mathbf{C}^\top \mathbf{C})^{-1} \mathbf{C}^\top \mathbf{Y}$$

of which $\widehat{\beta}_0$ is the first element. In more detail,

$$\widehat{\beta}_0 = \bar{Y} - \widehat{\beta}_1(\bar{C} - \mu_C) \quad (8.51)$$

where

$$\widehat{\beta}_1 = \frac{\sum_{i=1}^n (Y_i - \bar{Y})(C_i - \bar{C})}{\sum_{i=1}^n (C_i - \bar{C})^2}. \quad (8.52)$$

We refer to $\widehat{\beta}_0$ as a *control-variate estimator* of μ_Y and C as a *control*. Expression (8.51) provides additional intuition for this approach: The term $\widehat{\beta}_1(\bar{C} - \mu_C)$ adjusts the usual estimator \bar{Y} based on the difference between \bar{C} and its expected value μ_C .

Does $\widehat{\beta}_0$ have smaller variance than \bar{Y} ? Under very mild conditions $\widehat{\beta}_0$ satisfies the central limit theorem

$$\sqrt{n}(\widehat{\beta}_0 - \mu_Y) \xrightarrow{D} N(0, (1 - \rho^2)\sigma_Y^2) \quad (8.53)$$

as $n \rightarrow \infty$ (Nelson 1990). Thus, for large n

$$\text{Var}(\widehat{\beta}_0) \approx (1 - \rho^2) \frac{\sigma_Y^2}{n} \leq \frac{\sigma_Y^2}{n} = \text{Var}(\bar{Y}).$$

This shows that if we use the regression approach then the bivariate normal result holds in the limit even when (Y, C) are not bivariate normally distributed and β_1 must be estimated.

To provide a measure of estimation error on $\widehat{\beta}_0$ we can use the regression estimator of the variance-covariance matrix of $\widehat{\boldsymbol{\beta}}$:

$$\widehat{\boldsymbol{\Sigma}} = \widehat{\text{Var}}(\widehat{\boldsymbol{\beta}}) = (\mathbf{Y}^\top \mathbf{Y} - \widehat{\boldsymbol{\beta}}^\top \mathbf{C}^\top \mathbf{Y})(\mathbf{C}^\top \mathbf{C})^{-1} / (n - 2).$$

The (1,1) element is the estimated $\text{Var}(\widehat{\beta}_0)$, and an approximate confidence interval is

$$\widehat{\beta}_0 \pm t_{1-\alpha/2, n-2} \sqrt{\widehat{\boldsymbol{\Sigma}}_{11}}.$$

Since $\widehat{\Sigma}_{11}$ is the variance estimator for the intercept β_0 in regression it is generated by standard regression software. In detail, it is

$$\widehat{\Sigma}_{11} = \frac{\sum_{j=1}^n (Y_j - \widehat{\beta}_0 - \widehat{\beta}_1(C_j - \mu_C))^2}{n-2} \left(\frac{1}{n} + \frac{(\bar{C} - \mu_C)^2}{\sum_{j=1}^n (C_j - \bar{C})^2} \right).$$

This development makes clear that when we are looking for a control C two things are essential: The expected value of C must be known, and C should be strongly correlated with Y . Some possibilities from our standard examples follow:

8.5.1 M/G/1 Queue Control Variate

Recall that in the $M/G/1$ queueing simulation of Sect. 4.3 the output is the average of the last $m-d$ customer waiting times,

$$\bar{Y} = \frac{1}{m-d} \sum_{i=d+1}^m Y_i$$

where

$$\begin{aligned} Y_0 &= 0 & X_0 &= 0 \\ Y_i &= \max\{0, Y_{i-1} + X_{i-1} - A_i\}, & i &= 1, 2, \dots, m \end{aligned}$$

and we are interested in the steady-state expected waiting time $\mu = E(Y)$. One possible control is the average difference between the service time of customer $i-1$ and the interarrival time between customer $i-1$ and customer i :

$$\bar{C} = \frac{1}{m-d} \sum_{i=d+1}^m (X_{i-1} - A_i).$$

For the parameter settings in Sect. 4.3 $\mu_C = 1 - 0.8 = -0.2$.

When we make n replications we observe a pair of averages (\bar{Y}_j, \bar{C}_j) on each one. Therefore, the regression set up is

$$\begin{pmatrix} \bar{Y}_1 \\ \bar{Y}_2 \\ \vdots \\ \bar{Y}_n \end{pmatrix} = \begin{pmatrix} 1 & \bar{C}_1 - \mu_C \\ 1 & \bar{C}_2 - \mu_C \\ \vdots & \vdots \\ 1 & \bar{C}_n - \mu_C \end{pmatrix} \begin{pmatrix} \beta_0 \\ \beta_1 \end{pmatrix} + \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_n \end{pmatrix}.$$

Rerunning the example the control-variate estimator provides 95% confidence interval $[2.109, 2.187]$ as compared to $[2.101, 2.196]$ without it. The estimated correlation between the average waiting time \bar{Y} and \bar{C} is 0.648, a rather weak correlation which explains why there is little difference in the confidence-interval widths.

8.5.2 Stochastic Activity Network Control Variate

Recall that in the stochastic activity network of Sect. 4.4 the output is the time to complete the project

$$Y = \max\{X_1 + X_4, X_1 + X_3 + X_5, X_2 + X_5\}$$

and we are interested in $\theta = \Pr\{Y > t_p\}$. One possible control is $C = X_1 + X_3 + X_5$, the path with the longest expected value: Since X_1, X_3 and X_5 are independent, exponentially distributed random variables each with mean 1, C has an Erlang distribution with mean 3 and three phases. Using this fact we can show that

$$\theta_C = \Pr\{C > t_p\} = (1 + t_p + t_p^2/2)e^{-t_p}$$

which is easy to compute. The regression set up in this case is

$$\begin{pmatrix} I(Y_1 > t_p) \\ I(Y_2 > t_p) \\ \vdots \\ I(Y_n > t_p) \end{pmatrix} = \begin{pmatrix} 1 & I(C_1 > t_p) - \theta_C \\ 1 & I(C_2 > t_p) - \theta_C \\ \vdots & \vdots \\ 1 & I(C_n > t_p) - \theta_C \end{pmatrix} \begin{pmatrix} \beta_0 \\ \beta_1 \end{pmatrix} + \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_n \end{pmatrix}.$$

Since the output and the control are indicator (0 or 1) random variables, they clearly do not have a bivariate normal distribution; nevertheless, the central limit theorem result implies that we can attain a variance reduction when we have a large number of replications n .

Rerunning the example at $t_p = 5$, the control-variate estimator provides a 95% confidence interval $[0.165, 0.194]$ as compared to $[0.140, 0.186]$ without it. In this experiment $n = 1,000$. The estimated correlation between $I(Y > t_p)$ and $I(C > t_p)$ is 0.78, and there is a small, but noticeable reduction in confidence-interval width.

8.5.3 Asian Option Control Variate

This example is derived from Glasserman (2004, Chaps. 3 and 4).

Recall that in the Asian option example of Sect. 4.5 the goal was to estimate

$$v = E \left[e^{-rT} (\bar{X}(T) - K)^+ \right]$$

using the output $Y = e^{-rT} \max \left\{ 0, \widehat{\bar{X}}(T) - K \right\}$ where

$$\widehat{\bar{X}}(T) = \frac{1}{m} \sum_{i=1}^m X(i\Delta t).$$

A closely related problem is estimating

$$v_C = E \left[e^{-rT} (\tilde{X}(T) - K)^+ \right]$$

where

$$\tilde{X}(T) = \left(\prod_{i=1}^m X(i\Delta t) \right)^{1/m}$$

is the geometric average of the asset values. The value of an Asian option based on the geometric average is analytically solvable, so $C = e^{-rT} \max\{0, \tilde{X}(T) - K\}$ is a potential control. The formula for v_C is given in the Exercises.

Rerunning the example the control-variate estimator provides a 95% confidence interval [\$2.17, \$2.18] as compared to [\$2.10, \$2.29] without it. The estimated correlation between Y and C is 0.999, which explains the dramatic reduction in the confidence-interval width.

Appendix: Properties of Control-Variate Estimators

This appendix provides some small-sample (finite n) properties of control-variate estimators. All of these results assume that the output and the control are linearly related; we add assumptions as needed to obtain properties. These assumptions are stronger than required for the central limit theorem result (8.53), but lead to results that apply for finite n . Complete details can be found in Nelson (1990).

1. Under model (8.48), $E(\hat{\beta}_1) = \beta_1$.

Proof. For convenience, let $\mathbf{C}^\top = (C_1, C_2, \dots, C_n)$. Then

$$\begin{aligned} E(\hat{\beta}_1 | \mathbf{C} = \mathbf{c}) &= \frac{\sum E((Y_i - \bar{Y}) | \mathbf{C} = \mathbf{c})(c_i - \bar{c})}{\sum (c_i - \bar{c})^2} \\ &= \frac{\sum (\mu_Y + \beta_1(c_i - \mu_C) - \mu_Y - \beta_1(\bar{c} - \mu_C))(c_i - \bar{c})}{\sum (c_i - \bar{c})^2} \\ &= \beta_1. \end{aligned}$$

So that, by the double expectation theorem, $E(\hat{\beta}_1) = \beta_1$. □

2. Under model (8.48), $E(\hat{\beta}_0) = \mu_Y$.

Proof. $E(\hat{\beta}_0 | \mathbf{C} = \mathbf{c}) = E(\bar{Y} | \mathbf{C} = \mathbf{c}) - E(\hat{\beta}_1 | \mathbf{C} = \mathbf{c})(\bar{c} - \mu_C) = \mu_Y + \beta_1(\bar{c} - \mu_C) - \beta_1(\bar{c} - \mu_C) = \mu_Y$ using the result above. Since this is independent of \mathbf{c} the conclusion holds in general. □

Thus, the control-variate estimator is unbiased under the assumed linear model. When will it have smaller variance than \bar{Y} ? Notice that

$$\text{Var}(\hat{\beta}_0) = \text{Var} \left[E(\hat{\beta}_0 | \mathbf{C}) \right] + E \left[\text{Var}(\hat{\beta}_0 | \mathbf{C}) \right] = E \left[\text{Var}(\hat{\beta}_0 | \mathbf{C}) \right]$$

since $\text{Var}\left[\mathbb{E}(\widehat{\beta}_0|\mathbf{C})\right] = \text{Var}[\mu_Y] = 0$ from the proof of the result above. Still, this expression is not easy to evaluate in general. Under the additional assumption of constant conditional variance ($\text{Var}(Y|C) = \sigma^2$), the following result can be (tediously) derived:

3. Under model (8.48) and the assumption of constant conditional variance,

$$\text{Var}(\widehat{\beta}_0|\mathbf{C}) = \sigma^2 \left(\frac{1}{n} + \frac{(\bar{C} - \mu_C)^2}{\sum(C_i - \bar{C})^2} \right).$$

To complete the calculation we need to take the expectation of this result, which depends on the distribution of C . If (Y, C) are jointly normally distributed with correlation ρ , then we know that $\sigma^2 = (1 - \rho^2)\sigma_Y^2$ and it can be shown that

$$\mathbb{E} \left(\frac{(\bar{C} - \mu_C)^2}{\sum(C_i - \bar{C})^2} \right) = \frac{1}{n(n-3)}.$$

This leads to the following:

4. If (Y, C) are bivariate normal, then

$$\text{Var}(\widehat{\beta}_0) = \left(\frac{n-2}{n-3} \right) (1 - \rho^2) \frac{\sigma_Y^2}{n}.$$

Therefore, if $\rho^2 > 1/(n-2)$ then the control-variate estimator has smaller variance than the sample mean. Thus, provided the number of replications is not too small the control-variate estimator will have smaller variance even if the correlation is weak.

Exercises

1. Verify (8.6) and (8.7).
2. Show that for the AR(1) process the asymptotic MSE is

$$\text{MSE}(\bar{Y}(n, m, d)) \approx \frac{(y_0 - \mu)^2 \varphi^{2d+2}}{(m-d)^2(1-\varphi)^2} + \frac{\sigma^2}{n(m-d)(1-\varphi)^2}.$$

Then prove that for $\varphi > 0$ the asymptotic squared bias $(y_0 - \mu)^2 \varphi^{2d+2}/(m-d)^2(1-\varphi)^2$ is decreasing in d for m large enough. *Hint:* Take the derivative with respect to d .

3. Use the result in Exercise 2 to study the impact of dependence (as measured by φ) and initial conditions (as measured by $(y_0 - \mu)^2$) on the optimal deletion point d . A similar study was undertaken by Snell and Schruben (1985).
4. Using the result from Exercise 2, show that for the AR(1) process and m large

$$E(\text{MSER}(d)) \approx \frac{(1 - \phi)^2}{(1 - \phi^2)} \left(\frac{(y_0 - \mu)^2 \phi^{2d+2}}{(m - d)^2 (1 - \phi)^2} + \frac{\sigma^2}{n(m - d)(1 - \phi)^2} \right).$$

That is, the expected value of MSER is proportional to the asymptotic MSE.

5. As an aid to computing the MSER statistic, show that

$$\sum_{i=d+1}^m (Y_i - \bar{Y}(m, d))^2 = \sum_{i=d+1}^m Y_i^2 - \frac{1}{m-d} \left(\sum_{i=d+1}^m Y_i \right)^2.$$

6. For a covariance stationary process Y_1, Y_2, \dots, Y_m we know that the variance of the sample mean is

$$\text{Var}(\bar{Y}) = \frac{\sigma^2}{m} \left(1 + 2 \sum_{k=1}^{m-1} \left(1 - \frac{k}{m} \right) \rho_k \right)$$

where σ^2 is the marginal variance and ρ_k is the lag- k autocorrelation. If we have data, it would seem that we could estimate each term in this expression directly. What problems can you identify with this direct approach?

7. Remote Order Taking (ROT) is proposing to replace the current drive-through window ordering for a fast-food restaurant chain with the equivalent of a call center. ROT promises lower cost and faster response. Before agreeing to this, the corporate owner of the restaurant chain has asked for a proof-of-concept study using data from seven of their stores in Columbus, Indiana.

The typical drive-through window has a single traffic lane featuring a menu-board/order-entry station, a window where the person taking orders also collects payment, and a second window where the food is delivered to the customer. ROT is proposing a service that will eliminate the need for the first window and the corresponding employee by providing high-speed voice and data connections that will allow operators in Rapid City, South Dakota, to take orders and communicate them to the restaurant. The second window will then be used to collect payment and deliver the food to the customer.

This restaurant chain has high standards for customer service. When asked by ROT, they quantified this by saying that they want the average response time from when a drive-through customer's car triggers the sensor at the order board until that customer is greeted by the order taker to be 3 s or less, and the chance that any customer's response time is greater than 7 s be no more than 20%.

Whether or not ROT's proposal is worthwhile depends on whether the number of operators required to achieve these two goals is significantly less than one per store. The restaurant has provided data on the busiest 3-h period from the seven stores in the Columbus area to allow ROT to do a proof-of-concept study. You have been hired by ROT to simulate serving these seven Columbus stores to determine the minimum number of operators required to meet the service-level requirements.

Modeling the arrival process to each store as Poisson seems reasonable. From a prior study, data on the time required to take an order and the time for a car

Store number	Average number of customers (11 a.m. to 2 p.m.)
1	36
2	15
3	45
4	26
5	36
6	36
7	45

to clear the order board and the next car in line (if there is one) to pull up has been collected; that data can be found on the book website in `ROTDData.xls`. Clearly, arrival rates will vary by time of day. Since we only have data on the busiest period, treat this as a steady-state simulation where the long-run arrival rate is the same as the busy period; staffing that can keep up over the long-run at the heaviest load will certainly be adequate generally. Therefore, you will need to deal with initial-condition bias.

Notice that the entire food-preparation and food-delivery function is outside the scope of this project. You are only interested in order taking, and your goal is to investigate how many operators would be required to serve the seven stores.

8. Suppose that Y_1, Y_2, \dots, Y_m is a covariance stationary process with mean μ . Let

$$\hat{\sigma}^2 = \frac{1}{m} \sum_{i=1}^m (Y_i - \mu)^2$$

be the sample variance when the mean μ is known. Derive an expression for $E(\hat{\sigma}^2/m)$ and compare it to the expression in Exercise 6, the true variance of the sample mean. Is $\hat{\sigma}^2/m$ an unbiased estimator of $\text{Var}(\bar{Y}(m))$? Next do the same analysis for

$$S^2 = \frac{1}{m-1} \sum_{i=1}^m (Y_i - \bar{Y}(m))^2$$

the usual sample variance and compare $E(S^2/m)$ to the expression in Exercise 6.

9. When we only have a single replication $Y_1, Y_2, Y_3, \dots, Y_m$ of a steady-state simulation, an alternative to the MSER statistic for determining a deletion point is to plot the *cumulative average*

$$\bar{Y}(j) = \frac{1}{j} \sum_{i=1}^j Y_i$$

for $j = 1, 2, 3, \dots, m$, and truncate at the point when this plot becomes flat. To examine this idea, assume that the underlying output is an AR(1) process

$$Y_{i+1} = \mu + \phi(Y_i - \mu) + X_{i+1}$$

with $y_0 \neq \mu$, $\varphi > 0$ and X_1, X_2, \dots i.i.d. with mean 0 and variance σ^2 . Provide analysis that shows why this approach is conservative, meaning that it will tend to give a larger deletion point than necessary.

10. Recall the MA(1) surrogate process of Exercise 11 in Chap. 5:

$$Y_i = \mu + \theta X_{i-1} + X_i$$

with $X_0 = x_0$ a fixed constant, and X_1, X_2, \dots i.i.d. with mean 0 and variance σ^2 , and $|\theta| < 1$. That exercise asked you to derive the asymptotic MSE of the sample mean $\bar{Y}(m) = \sum_{i=1}^m Y_i/m$ with no data deletion. Extend your result to obtain the asymptotic MSE for $d \geq 1$. If m is fixed, could it ever be optimal (minimize asymptotic MSE) to have $d > 0$, and if so what is the optimal d ?

11. Recall the MA(1) surrogate process of Exercise 11 in Chap. 5:

$$Y_i = \mu + \theta X_{i-1} + X_i$$

with $X_0 = x_0$ a fixed constant, X_1, X_2, \dots i.i.d. with mean 0 and variance σ^2 , and $|\theta| < 1$. For output data from an MA(1) process, compute the expected value of the MSER(d) statistic as a function of $d = 0, 1, 2, \dots, m - 1$. After how many observations do we expect deletion?

12. In Exercises 2 and 10 you derived the asymptotic MSE for AR(1) and MA(1) surrogate processes. Using these expressions you can approximate the MSE as a function of the initial conditions, run length m , deletion point d and values for the process parameters (φ, θ , etc.). Perform an empirical study to see how effective the minimum MSER statistic is at minimizing the MSE. In other words, simulate various cases of the AR(1) and MA(1) processes, estimate a deletion point using the minimum MSER rule, and evaluate how good the chosen deletion point is using the known asymptotic MSE.
13. Recall the $M/G/1$ simulation of Sect. 4.3, and specifically the steady-state waiting time Y . Using a single replication and batching, estimate and provide a 90% confidence interval for the marginal variance and the 0.8 quantile of Y .
14. If you could initialize a steady-state simulation in “steady-state conditions,” then there would be no initial-condition bias. Since this can rarely be done, it has been suggested that two replications of the simulation could be made such that the first is initialized in overloaded conditions, while the second is initialized in underloaded conditions, and then the two replications could be averaged to reduce the bias. For instance, the first replication of a queueing simulation could be initialized with many customers in the system, while the second replication could be initialized empty and idle. Assuming that the underlying output is an AR(1) process

$$Y_{i+1} = \mu + \varphi(Y_i - \mu) + X_{i+1}$$

with y_0 being a constant representing the initial conditions and X_i i.i.d. $(0, \sigma^2)$, show mathematically whether or not this idea always reduces bias relative to picking one fixed initial condition.

15. Recall the paired- t confidence interval (8.26) for the difference $\theta(\mathbf{x}_1) - \theta(\mathbf{x}_2)$. Show that

$$S_D^2 = S^2(\mathbf{x}_1) + S^2(\mathbf{x}_2) - 2S(\mathbf{x}_1, \mathbf{x}_2)$$

where $S^2(\mathbf{x}_1)$ and $S^2(\mathbf{x}_2)$ are the marginal sample variances, and

$$S(\mathbf{x}_1, \mathbf{x}_2) = \frac{1}{n-1} \sum_{j=1}^n (Y_j(\mathbf{x}_1) - \bar{Y}(\mathbf{x}_1)) (Y_j(\mathbf{x}_2) - \bar{Y}(\mathbf{x}_2))$$

is the sample covariance. This shows how the paired- t interval captures the effect of common random numbers.

16. Consider the following five scenarios for the mean activity times in the SAN example. Make 100 replications of each one and apply subset selection with 95% confidence to see which if any can be eliminated from being the best in terms of the smallest expected value of the time to complete the project. Now repeat with 200 replications of each.

	x_1	x_2	x_3	x_4	x_5
x_{i1}	0.5	1	1	0.3	1
x_{i2}	1	0.5	1	1	1
x_{i3}	1	1	0.5	1	1
x_{i4}	1	1	1	0.4	1
x_{i5}	1	1	1	1	0.5

17. For the same scenarios as in Exercise 16, apply the selection-of-the-best procedure to find the scenario with the smallest expected completion time with 95% confidence. Use $n_0 = 10$ and $\delta = 0.1$. Run the experiment with and without using common random numbers and comment on the number of replications required to reach a decision in each case.
18. Implement all of the gradient estimators described in this chapter for the SAN example. Using $\tau_j = c_j = 1$ and $\ell_j = 0.5$ for $j = 1, 2, 3, 4, 5$, and $b = 1$, select several feasible settings for \mathbf{x} and estimate the gradient using each method. Compare the methods in terms of the variability of the gradient estimates. Use $\Delta x = 0.1$ in FD.
19. When applying stochastic approximation to the SAN optimization, suppose that at iteration i we find $\sum_{j=1}^5 c_j(\tau_j - x_{i+1,j}) > b$. First derive \mathbf{x}'_{i+1} , the perpendicular projection of \mathbf{x}_{i+1} onto the plane $\sum_{j=1}^5 c_j(\tau_j - x_j) = b$. Then provide a refinement that insures that \mathbf{x}'_{i+1} also satisfies the lower bound constraints $x'_{i+1,j} \geq \ell_j$, $j = 1, 2, 3, 4, 5$; this need not be a point that is on the plane, but it should be feasible.
20. Using each of the gradient estimators described in this chapter, implement a stochastic approximation search for the optimal mean activity times for the SAN example using the parameters in Exercise 18. Suggested settings for this problem are $a_i = 1/i$, starting scenario $\mathbf{x}^\top = (1, 1, 1, 1, 1)$, and at least 50 replications per gradient estimate, but you should experiment with these values.

21. Using linear programming software, solve the SAA problem (8.39) using $\tau_j = c_j = 1$ and $\ell_j = 0.5$ for $j = 1, 2, 3, 4, 5$, and $b = 1$. Use $n = 10$ replications if you have to set it up manually, and $n = 10, 100, 1,000$ if you can automate the set up.
22. Work out what the variance-decomposition result (8.47) implies when (a) Y and C are independent, and when (b) $Y = C$.
23. Show that

$$\beta_1^* = \frac{\text{Cov}(Y, C)}{\text{Var}(C)} = \frac{\rho \sigma_Y \sigma_C}{\sigma_C^2}$$

minimizes the variance of the random variable $\mu_Y + \beta_1^*(C - \mu_C)$. This provides another justification for the control-variate estimator.

24. For the SAN example, use a control-variate estimator to estimate $\mu_Y = E(Y)$, the expected time to complete the project. Compare it to just using \bar{Y} .
25. The control-variate estimator can employ more than one control, analogous to using more than one explanatory variable in linear regression. Try this idea on the $M/G/1$ simulation forming one control-variate based on the interarrival times, and another based on the service times.
26. The control-variate estimator can employ more than one control, analogous to using more than one explanatory variable in linear regression. Try this idea on the SAN simulation forming controls for each of the three paths through the network.
27. Specializing results in Glasserman (2004, Chap. 3) to our Asian option example (where the value of the asset is recorded at m equally spaced time intervals of length Δt), it can be shown that

$$v_C = e^{-\delta T'} X(0) \Phi(d) - e^{-rT'} K \Phi(d - \bar{\sigma} \sqrt{T'})$$

where $T' = (m + 1)\Delta t/2$, $\bar{\sigma}^2 = (2m + 1)\sigma^2/(3m)$, $\delta = (\sigma^2 - \bar{\sigma}^2)/2$ and

$$d = \frac{\ln(X(0)/K) + (r - \delta + \bar{\sigma}^2/2)T'}{\bar{\sigma} \sqrt{T'}}$$

Use this result to implement the control-variate estimator for the Asian call option. Test the effect of different numbers of discretization steps $m = 8, 16, 32, 64, 128$.

Chapter 9

Simulation for Research

This book is about simulation modeling, programming and experimentation for the purpose of systems analysis. However, stochastic simulation is also a tool that can be used to support basic research in domains such as simulation, optimization, queueing, financial engineering, production planning and logistics. In this chapter we cite some papers that demonstrate effective application of simulation in research and use them to highlight general principles and practices. We start with two important distinctions.

The first distinction is between a *practitioner's experiment* and a *researcher's experiment*. This book is written primarily from the point of view of the practitioner's experiment. The practitioner has a real problem to solve and has some limit on how much time or effort they can expend to solve it. The practitioner will try to build an appropriate simulation model or models, will perform the simulation experiment using the best methods they know or have access to, and will use the results to solve their problem. If they learned simulation from this (or many other) books then they will also try to assess the quality of their answer (e.g., via a confidence interval) or they will use a procedure that is designed to deliver a certain level of quality (e.g., a ranking and selection procedure with a guaranteed probability of correct selection). They will, however, only have a vague sense as to whether they got the "right answer" to their problem based on what happens when they apply their simulation results in the real world.

The researcher's experiment, on the other hand, is driven by an abstract research question whose answer may be useful in practice, but is not the answer to a specific practical problem. Because it is not directly tied to the real world, which is messy, a research question can be formulated and answered very precisely. The researcher's experiment may, and often will, involve solving the same problem repeatedly for simulated instances that differ through some controllable experimental factors and in the random numbers used in the simulation. For instance, Sect. 9.1 below describes generating random optimization problems of a particular type and solving each instance with the same algorithm to assess the algorithm's performance over a space of problems.

Unlike the practitioner's experiment, the researcher's experiment frequently involves simulating scenarios for which the answer is known in advance; knowing the answer facilitates a better evaluation. For example, Sect. 9.3 describes an experimental study of new ranking and selection procedures. Ranking and selection procedures are supposed to discover the best scenario with a guaranteed probability of correct selection. The achieved probability of correct selection for a new procedure can be estimated by applying it repeatedly (with different random numbers) to simulated scenarios for which the true best scenario is already known. This contrasts with the practitioner who will apply ranking and selection once on a collection of scenarios for which the best is not known and then actually implement their selection.

A second distinction is between simulation *experiments* that support research as opposed to *illustrations* or examples. What is the difference? A successful experiment allows the researcher to make statements about cases, scenarios or problems that they *did not* try based on cases, scenarios or problems that they *did* try. To accomplish this an experiment must be *designed* so as to provide inference about a well-defined space of cases, scenarios or problems. An illustration, on the other hand, is a specific case, scenario or problem that helps the reader understand how a method is implemented and how the results should be interpreted.

This book is full of illustrations, starting with those in Chap. 3 that were used to demonstrate simulation methods throughout the book. These illustrations were (we hope) useful for learning and understanding the methods, but by themselves they do not prove that the methods work more generally.

In the following sections we provide some guidelines for generating random test problems, conducting robustness studies, linking experimental factors and measuring and controlling error in a research simulation experiment. In addition to the research papers cited, this chapter also draws on Sect. 3 of Goldsman et al. (2002) which discusses evaluation of simulation output analysis procedures.

9.1 Random Test Problems

Simulation is often used to generate random test problems to support saying something general about the performance of a solution method. Consider the following example:

The two-dimensional knapsack problem is a deterministic (and difficult) combinatorial optimization problem of the form

$$\begin{aligned} \max \sum_{j=1}^n c_j x_j & & (9.1) \\ \sum_{j=1}^n a_{ij} x_j \leq b_i, & \quad i = 1, 2 \\ x_j \in \{0, 1\}, & \quad j = 1, 2, \dots, n \end{aligned}$$

where $c_i > 0$ and $a_{ij} \geq 0$. Think of c_j as the value of selecting the j th item ($x_j = 1$), while a_{1j} and a_{2j} are two types of “costs” for selecting it (e.g., weight and price), and the totals of these “costs” are constrained by b_1 and b_2 , respectively.

Hill and Reilly (2000) compared the performance (speed and quality of solution) of a heuristic and a convergent algorithm for solving (9.1) by applying them to a large number of randomly generated test problems. A test problem is specified by providing values for c_j, a_{ij}, b_j and n . Hill and Reilly (2000) fixed the number of items at $n = 100$ and set b_i to be a specified fraction of $\sum_{j=1}^n a_{ij}$. This left them with the need to decide how to randomly generate values of c_i and a_{ij} to draw meaningful comparisons.

A prescription that seems reasonable is to select uniform distributions for c_j, a_{1j} and a_{2j} , and then to randomly and independently generate the necessary values. We call these “completely random test problems.” Assuming that we are comfortable setting the ranges on these uniform distributions, then the space specified by the ranges will be randomly covered. But is that enough?

Hill and Reilly (2000) showed conclusively that the answer is “no” by demonstrating that *correlation* across the coefficients (c_j, a_{1j}, a_{2j}) can have a dramatic effect on optimization performance, and these correlations have a practical meaning. For instance, a strong positive correlation between a_{ij} and c_j means that if the j th item is valuable, then it also tends to be costly; and this makes the problem hard because desirable items tend to break the budget. Similarly, strong negative correlation between a_{1j} and a_{2j} means that an item that is inexpensive relative to the budget b_1 tends to be expensive relative to budget b_2 ; this also makes the problem difficult. Further, the correlation structure affects the heuristic and convergent algorithm differently. While some problems that exhibit these features will be generated, by chance, for completely random test problems, they will be rare and their effect will be averaged over all of the (mostly) low-correlation test problems. Therefore, this important, and practically justifiable, feature will not be identified.

Hill and Reilly (2000) illustrates our first principle of simulation for research:

Research principle 1. *Completely random cases, scenarios or problems are not necessarily relevant ones. Instead, generate test cases, scenarios or problems that have features that are representative of the space of interest.*

Hill and Reilly (2000) carefully controlled the strength of the correlations in their random test problems via a designed experiment. This allowed them to make much stronger statements about how the solutions methods will behave in practice. Variate-generation methods such as NORTA (Sect. 6.3.2) can be used to induce such correlation.

This example also illustrates an obvious, but important second principle:

Research principle 2. *If you are comparing things, use common random numbers.*

Hill and Reilly (2000) could have generated the test problems to evaluate the heuristic independently from those used to evaluate the convergent algorithm. That is, they could have chosen the correlation settings at which they wanted to do experiments, but used different random numbers to actually generate the coefficients for

the problems solved by the heuristic and for the problems solved by the algorithm. The comparisons would still have been with respect to the same problem space, but not with respect to precisely the same knapsack problems. If they had used independent test sets, then some of the observed differences in performance would have been due to solving different test problems; this source of variation is eliminated by using common random numbers since the problem sets will be identical.

9.2 Robustness Studies

Simulation experiments are also used to assess the robustness of a research result. By “robustness” we mean the ability of the procedure, method or concept to perform as expected across a wide range of conditions.

Controlled experimentation is a topic addressed by the statistical design of experiments; see for instance Montgomery (2009). Experiment design starts by identifying the factors that might affect the performance of a product, process or treatment and then varies these factors in a systematic fashion. Careful identification of the appropriate factors, both favorable and unfavorable, is critical to a study of robustness, as illustrated in the following example.

Iravani et al. (2005) proposed indices for measuring the structural flexibility of a manufacturing or service operation. Flexibility in their context is the capability of the operation to satisfy multiple types of demand with available resources in the face of changing demand and resource capacity. The “structural” in structural flexibility refers to the choice of how many and which types of demand that each resource can satisfy. In a manufacturing setting where there are N factories and K types of products (e.g., car models), structural flexibility depends on which of these products each factory can produce. In a service setting like a call center where there are N agents and K types of calls, structural flexibility depends on which types of calls each agent is cross-trained to handle. Iravani et al. (2005) proposed an easy-to-compute structural flexibility index for strategic-level comparison of competing system designs, an index that should predict which system design would be most productive in a changing and unpredictable real-world environment.

In brief, Iravani et al. (2005) represented a system design as a graph from resources $\{S_1, S_2, \dots, S_K\}$ to demand types $\{D_1, D_2, \dots, D_N\}$ with an arc from $S_k \rightarrow D_j$ if resource k can satisfy some demand of type j . Their two structural flexibility (SF) indices are functions of the number of nonoverlapping paths through this network by which excess capacity to satisfy demand type i can be redirected to satisfy demand of type j , and also the number of different resources that can satisfy demand of type j , for $i, j = 1, 2, \dots, K$. The research question is, how robust is this simple measure for ranking the performance of system designs relative to the types of real-world systems to which it might be applied?

To answer this question Iravani et al. (2005) employed a designed simulation experiment. The key to designing an experiment that convincingly establishes ro-

business is that it vary factors that are *not* inputs to the SF measure, but do represent what occurs in the real world. These included the following:

Physical flow in the system: An open, parallel queueing environment (such as a call center) and a closed serial queueing environment (such as a production line operating with a constant work in process or WIP) were chosen.

How close the system is to capacity: Congestion levels were controlled in the open queueing system by setting the demand arrival rate to achieve specific levels of utilization, and in the closed system by setting the WIP to achieve different levels of jobs/worker.

Uncertainty in the environment: Short-term fluctuations and long-term changes were considered:

Short-term variation: Interarrival times and processing times were modeled as having gamma distributions, with different levels of variability attained by setting their coefficient of variation (standard deviation divided by the mean).

Long-term variation: Shocks randomly inject a persistent change in demand. Shocks that increase one type of demand at the expense of another, and shocks that simply increase one type of demand were included.

The simulation experiment was conducted on pairs of system designs to see if the one with the higher SF index (which is computable without simulation) was also the one with the higher productivity (as estimated via simulation).

This paper illustrates another research principle:

Research principle 3. *Run a designed experiment that identifies and controls factors that might influence the outcome, both favorably and unfavorably, and that are actually encountered in practice.*

9.3 Linking Experimental Factors

In physical experiments the operable range for each experimental factor, such as temperature, pressure or dosage, can often be set based on the nature of the system. Unfortunately, this may not be the case when evaluating, say, a new simulation output analysis method that is supposed to apply to any problem in a large and diverse class. Therefore, rather than independently specifying a range for each factor, it may be better to link the ranges of the factors in meaningful ways. Here we provide one illustration.

Ranking and selection procedures are simulation optimization methods whose objective is to discover the best simulated system, in terms of largest or smallest true expected performance, with some guaranteed probability of correct selection. See Sect. 8.4.2 for background on ranking and selection. Nelson et al. (2001) proposed and then empirically evaluated ranking and selection procedures created for simulation optimization problems where the number of alternative scenarios K is moderate (up to 500). Their interest was primarily in computational efficiency of the procedures as measured by the number of simulation replications required to deliver the guarantee.

A number of factors in addition to K can affect the efficiency of a ranking and selection procedure; Nelson et al. (2001) focused on four: The values of the true means, $\mu_1, \mu_2, \dots, \mu_K$, the output variances of each scenario, $\sigma_1^2, \sigma_2^2, \dots, \sigma_K^2$, the initial (“first stage”) number of replications obtained from each scenario n_0 , and the size of the indifference-zone parameter δ . Recall that ranking and selection procedures are often sequential and therefore have multiple stages, and that δ is the smallest practically significant difference in mean performance that the user is interested in detecting. In their experiments, Nelson et al. (2001) linked these factors to allow general conclusions to be drawn.

Notice that a ranking and selection problem will be computationally easy if the mean values are widely separated, and if the variances of the outputs are small while δ is large relative to the differences in the means. On the other hand, if the means are close, the variances large, and the difference we care about is small, then a great deal of simulation effort will be required to guarantee finding the best. These insights are useful for linking the factors.

Nelson et al. (2001) let $\mu_1 = \delta$ be the largest mean (bigger was better in their paper) and anchored the experiment design to this value. They considered a number of configurations of the means, including:

Slippage configuration: $\mu_1 = \delta$, or a multiple of δ , and $\mu_2 = \mu_3 = \dots = \mu_K = 0$.

This is a difficult configuration since all inferior scenarios are tied for second best.

Monotone decreasing means: $\mu_1 = \delta$, and $\mu_i = \mu_1 - (i - 1)\delta/\tau$, $i = 2, 3, \dots, K$, with $\tau = 1, 2, 3$. These are easier configurations since many scenarios are substantially inferior to the best.

They then set $\delta = d\sigma_1/\sqrt{n_0}$, a multiple of the standard deviation of the first-stage sample mean of the best system, varying $d = 1/2, 1, 2$. This ties the separation between the best and the rest to the precision with which we can estimate the mean of the best.

Finally, for the variances of the systems, they considered situations in which the variance of the best system is higher or lower than the inferior systems, as follows:

Common variance configuration: $\sigma_2^2 = \sigma_3^2 = \dots = \sigma_K^2 = \sigma^2$ and $\sigma_1^2 = \rho\sigma^2$, with $\rho = 1/2, 2$.

Unequal variance configuration: Variances increase as the mean decreases, $\sigma_i^2 = |\mu_i - \delta| + 1$, and variances decrease as the mean decreases, $\sigma_i^2 = 1/(|\mu_i - \delta| + 1)$, $i = 1, 2, \dots, K$.

Notice that the number of factors has now been effectively reduced to three: The number of scenarios K and first-stage sample size n_0 , both factors for which a reasonable range of operability can be determined, and σ_1^2 , the variance of the output from the best system, which can be set to any fixed value (say 1) because all of the other factors are relative to it. Thus, conclusions from the empirical study are with respect to the relative relationships among the means, variance and indifference zone of a simulation optimization problem, not their actual values.

Research principle 4. *When there are no natural ranges of operability for your factors, link them so that they are large or small relative to each other.*

Nelson et al. (2001) generated simulation output data directly from normal and lognormal distributions; this made it easy to control the experimental factors described above, and to know which scenario is actually the best. Controlling experimental factors is often easier with surrogate models than with realistic discrete-event simulations; surrogate models have characteristics like realistic simulations, but are simple, tractable and controllable. In this book we have often used the AR(1) as a surrogate for output from a steady-state simulation.

While well-chosen surrogate models share characteristics with real simulations, we cannot be certain that they include all characteristics that appear in, say, a complex supply chain simulation, and that might matter. For this reason, running experiments with models that are less controllable, but a step closer to realism, is often appropriate. Tractable queueing models (e.g., $M/M/\infty$, $M/G/1$), simple inventory models (e.g., (s, S) inventory model with Poisson demand) and simple financial options (e.g., European call or put) are common choices.

Research principle 5. *Since it may not be possible to anticipate all important factors, include some realistic examples along with the controllable surrogate models.*

9.4 Controlling Error in Research Simulations

Recall the Pollaczek–Khinchine formula (Eq. (3.4)) which gives the steady-state expected waiting time in an $M/G/1$ queue. A related formula provides the steady-state expected number of customers in the queue:

$$q = \frac{\lambda^2(\sigma^2 + \tau^2)}{2(1 - \lambda\tau)} = \frac{\rho^2(c^2 + 1)}{2(1 - \rho)}$$

where λ is the arrival rate of the Poisson arrival process and (τ, σ^2) are the mean and variance of the service-time distribution; or equivalently $\rho = \lambda\tau$ is the traffic intensity (a measure of system load), and $c = \sigma/\tau$ is the coefficient of variation of the service time. Typically we use this formula to predict the expected number in the queue given estimates for λ , τ and σ^2 . However, it is clear that if instead we had estimates of the service-time parameters (τ, σ^2) and the average number of customers in the queue q then we could use this formula to solve for what the arrival rate λ must have been. Having this estimated arrival rate we could then predict what the expected number in the queue would be if we changed the service process (e.g., lowered the mean service time, τ).

Whitt (1981) considered a more general version of this situation: Suppose that the congestion in a real-world queueing system can be observed, as well as the queue's service process, but the arrival process—which may be a very general stationary arrival process—is not observed or easily characterized. We want to know what the

expected number in the system would be if the same arrival process encountered an altered service process, new equipment or people, for instance.

Whitt's goal was to use the available information to construct a robust, simple, and tractable renewal-process approximation of the unknown arrival process that could then be the input to a queueing model. "Robust" in this context means that the queueing model will provide an accurate approximation of the steady-state congestion under the new service process; "simple" means that we only need two moments (arrival rate and the coefficient of variation of the time between arrivals) to fit the approximating renewal process; while "tractable" means that using this approximating renewal process we can compute the steady-state expected number in the system using something like a Pollaczek–Khinchine formula.

For instance, suppose that the observable queueing system is a $G/M/1$: general stationary arrival process G (which may include dependence), exponential service-time distribution and one server. The service rate and some measures of congestion for this system are observable. The altered system will be a $G/M'/1$, where M' is an exponentially distributed service process with a different service rate. The new system will be approximated by a $GI/M'/1$ queue; GI denotes "general, independent" which means a renewal arrival process. The objective is to have GI approximate G in such a way that the $GI/M'/1$ gives accurate congestion predictions for the $G/M'/1$. *A key point of the paper is that the quality of the arrival-process approximation is judged by how accurately the queueing model predicts steady-state congestion.*

To provide a fair evaluation, Whitt (1981) considered queues with arrival processes that were not of the same type as his approximating renewal processes. For instance, the Poisson arrival process was one of his approximations. Therefore, the $M/M/1$ queue was not an interesting test case because the true arrival process and the approximating arrival process are of the same type (Poisson). Instead, his test cases had arrival processes for which the steady-state expected number in system is not known.

How do you evaluate the accuracy of an approximation when you do not know the true value that it is approximating? Whitt (1981) used simulation to estimate the true expected number in the queue for his test cases; simulation is another form of approximation. To be valid, the error in the simulation estimate has to be small enough that it can be considered insignificant. A feature of simulation is that we can quantify the error in our estimates, and we can drive this error to 0 by increasing the run length or number of replications. Section 8.1 described methods for controlling the error in this way.

Research principle 6. *If you reduce the standard error or confidence interval width enough, then the simulation estimate is effectively the same as the true value.*

To be more concrete, suppose that Whitt's approximation is considered good enough if it has no more than 10 % error relative to the true value:

$$\frac{|q_{\text{approx}} - q_{\text{true}}|}{q_{\text{true}}} \stackrel{?}{\leq} 0.1$$

Instead of the true value, we have \hat{q}_{true} , a simulation estimate of it; \hat{q}_{true} has estimated error of, say, $\pm H$, which depends on the run length or number of replications. This error must be made small enough that we can conclude that the approximation is, or is not, within 10 % relative error even considering the estimation error. Stated differently, the ratio above should be satisfied, or not satisfied, for every possible value of q_{true} such that $\hat{q}_{\text{true}} - H \leq q_{\text{true}} \leq \hat{q}_{\text{true}} + H$. In Whitt (1981) this meant that different run lengths were required depending on the traffic intensity in the simulated queue.

We next consider a different aspect of controlling error, and a different paper. Sections 9.1–9.3 described choosing and controlling the important factors in a research simulation. This is not enough, however, when empirical evaluation of performance is estimated through repeated trials, as is common in the study of new statistical procedures. The illustration we will use is the evaluation of control-variate estimators, which were described in Sect. 8.5.

In brief, a control-variate estimator $\hat{\beta}_0$ is an alternative to the sample mean \bar{Y} for estimating the expected value $\mu_Y = E(Y)$. As the number of replications goes to infinity, control-variate estimators are asymptotically consistent for μ_Y and have smaller variance than \bar{Y} ; but in finite samples they may be biased and could have larger variance. Suppose we wanted to study how control-variate estimators behave when the sample size is small. We might be interested in properties such as bias, $E(\hat{\beta}_0 - \mu_Y)$, and variance, $\text{Var}(\hat{\beta}_0)$. Since there is an associated confidence interval for μ_Y , $\hat{\beta}_0 \pm t_{1-\alpha/2, n-2} \sqrt{\hat{\Sigma}_{11}}$, we might also be interested in its coverage probability and how well $\hat{\Sigma}_{11}$ works as an estimator of $\text{Var}(\hat{\beta}_0)$. Experiments such as the ones described below were used to study control-variate estimators in Nelson (1990).

Listing the factors that might affect the performance of control-variate estimators would provide a number of instances to simulate (see Nelson 1990). Most likely we would select instances for which μ_Y is known. For a particular instance, simulating it one time yields one set of estimates $\hat{\beta}_0^{(1)}$, $\hat{\Sigma}_{11}^{(1)}$, and $\hat{\beta}_0^{(1)} \pm t_{1-\alpha/2, n-2} \sqrt{\hat{\Sigma}_{11}^{(1)}}$, where we use a superscript (i) to indicate the i th trial. To estimate bias and variance, we need multiple trials or “macroreplications” of the same instance, $\hat{\beta}_0^{(1)}, \hat{\beta}_0^{(2)}, \dots, \hat{\beta}_0^{(m)}$. These will be i.i.d. because they are simulations of the same problem instance, but with different random numbers on each macroreplication. The bias, variance and coverage of the control-variate estimator for this instance can be estimated, respectively, by

$$\begin{aligned} \hat{b} &= \frac{1}{m} \sum_{i=1}^m \hat{\beta}_0^{(i)} - \mu_Y = \bar{\beta}_0 - \mu_Y \\ S_{\hat{\beta}_0}^2 &= \frac{1}{m-1} \sum_{i=1}^m \left(\hat{\beta}_0^{(i)} - \bar{\beta}_0 \right)^2 \\ \hat{p} &= \frac{1}{m} \sum_{i=1}^m I \left\{ \mu_Y \in \hat{\beta}_0^{(i)} \pm t_{1-\alpha/2, n-2} \sqrt{\hat{\Sigma}_{11}^{(i)}} \right\}. \end{aligned}$$

Each of these is an estimate, and therefore subject to error, but they are also averages of i.i.d. observations so we know how to measure and control the error to be small enough to support the conclusions we want to reach; m should not be chosen arbitrarily.

For instance, suppose that we set $1 - \alpha = 0.95$, meaning 95 % confidence intervals. Even if the confidence interval has the desired coverage, the standard error of \hat{p} is $\sqrt{(0.95)(0.05)/m}$. If $m = 30$ then two times this standard error is ≈ 0.08 , which means we cannot really distinguish coverage of 0.95 from 0.87. To get coverage estimates with approximately two decimal places of precision we need m to satisfy

$$2\sqrt{\frac{(0.95)(0.05)}{m}} < 0.01$$

or $m \approx 1900$.

What about assessing the quality of $\hat{\Sigma}_{11}$? Similar to the situation in Whitt (1981), the true $\text{Var}(\hat{\beta}_0)$ is not known. However, $S_{\hat{\beta}_0}^2$ is an unbiased estimator of it, since it is a direct estimator based on i.i.d. observations of $\hat{\beta}_0$. Therefore, an estimator of the bias of $\hat{\Sigma}_{11}$ is

$$\frac{1}{m} \sum_{i=1}^m \hat{\Sigma}_{11}^{(i)} - S_{\hat{\beta}_0}^2.$$

The analysis described above is embedded within the larger experiment that varies the factors that affect control-variate estimators; that is, it is an experiment nested inside an experiment. While the practitioner faces one problem instance, and solves it $m = 1$ times, the researcher creates (systematically or randomly) many instances, and simulates each to an acceptable level of error to draw valid conclusions. This gives the final research principle is, which is

Research principle 7. *Measure and control the error in your research experiment using nested simulations.*

Appendix A

VBASim

A.1 Core Subs

```
' THIS IS VBASIM 7/26/2009

' VBAsim is a minimal collection of Subs, Functions and Class Modules to support
' discrete-event simulation programming in VBA for Excel.

' Set up event handling

Public Clock As Double           'simulation global clock
Public Calendar As New EventCalendar 'event calendar

' Set up Collections to be reinitialized between replications

Public TheCTStats As New Collection ' continuous-time statistics
Public TheDTStats As New Collection ' discrete-time statistics
Public TheQueues As New Collection ' queues
Public TheResources As New Collection ' resources
```

```
Public Sub VBASimInit()  
  ' Sub to initialize VBASim  
  ' Typically called before the first replication and between replications  
  
  Dim Q As FIFOQueue  
  Dim CT As CTStat  
  Dim DT As DTStat  
  Dim EV As EventNotice  
  Dim En As Entity  
  Dim Re As Resource  
  
  ' Reset clock and empty event calendar  
  Clock = 0  
  Do While Calendar.N > 0  
    Set EV = Calendar.Remove  
    Set EV = Nothing  
  Loop  
  
  ' Empty queues  
  For Each Q In TheQueues  
    Do While Q.NumQueue > 0  
      Set En = Q.Remove  
      Set En = Nothing  
    Loop  
  Next  
  
  ' Clear statistics  
  
  For Each CT In TheCTStats  
    CT.Clear  
  Next  
  For Each DT In TheDTStats  
    DT.Clear  
  Next  
  
  ' Reinitialize Resources to idle  
  
  For Each Re In TheResources  
    Re.Busy = 0  
  Next  
  
End Sub
```

```
Public Sub Schedule(EventType As String, EventTime As Double)
' Schedule future events of EventType to occur at time Clock + EventTime

    Dim addedEvent As New EventNotice
    addedEvent.EventType = EventType
    addedEvent.EventTime = Clock + EventTime

' insert problem specific event notice attributes here

' schedule
    Calendar.Schedule addedEvent
    Set addedEvent = Nothing
End Sub
```

```
Public Sub Report(Output As Variant, WhichSheet As String, _
                Row As Integer, Column As Integer)
' Basic report writing sub to put Output on worksheet
' WhichSheet(Row, Column)

Worksheets(WhichSheet).Cells(Row, Column) = Output

End Sub
```

```
Public Sub SchedulePlus(EventType As String, EventTime As Double, _
                    TheObject As Object)
' Schedule future events of EventType to occur at time Clock + EventTime
' and pass TheObject

    Dim addedEvent As New EventNotice
    addedEvent.EventType = EventType
    addedEvent.EventTime = Clock + EventTime
    Set addedEvent.WhichObject = TheObject

' insert problem specific event notice attributes here

' schedule
    Calendar.Schedule addedEvent
    Set addedEvent = Nothing

End Sub
```

```

Public Sub ClearStats()
' Clear statistics in TheDTStats and TheCTStats

Dim CT As CTStat
Dim DT As DTStat

    For Each CT In TheCTStats
        CT.Clear
    Next
    For Each DT In TheDTStats
        DT.Clear
    Next

End Sub

```

A.2 Random-Number Generation

```

' This random number generator and random-variate generation functions are
' VBA translations of the C programs found in
' Law, A. M. and Kelton, W. D., ''Simulation Modeling and Analysis'',
' Singapore: The McGraw-Hill Book Co, pp. 430--431.

```

```

Option Explicit

Const MODLUS = 2147483647
Const MULT1 = 24112
Const MULT2 = 26143

' Define Static variable
Dim zrng() As Long

Public Static Sub InitializeRNSeed()

    ReDim zrng(1 To 100) As Long

    zrng(1) = 1973272912
    zrng(2) = 281629770
    zrng(3) = 20006270
    zrng(4) = 1280689831
    zrng(5) = 2096730329
    zrng(6) = 1933576050
    zrng(7) = 913566091
    zrng(8) = 246780520
    zrng(9) = 1363774876
    zrng(10) = 604901985
    zrng(11) = 1511192140
    zrng(12) = 1259851944
    zrng(13) = 824064364

```

zrng(14) = 150493284
zrng(15) = 242708531
zrng(16) = 75253171
zrng(17) = 1964472944
zrng(18) = 1202299975
zrng(19) = 233217322
zrng(20) = 1911216000
zrng(21) = 726370533
zrng(22) = 403498145
zrng(23) = 993232223
zrng(24) = 1103205531
zrng(25) = 762430696
zrng(26) = 1922803170
zrng(27) = 1385516923
zrng(28) = 76271663
zrng(29) = 413682397
zrng(30) = 726466604
zrng(31) = 336157058
zrng(32) = 1432650381
zrng(33) = 1120463904
zrng(34) = 595778810
zrng(35) = 877722890
zrng(36) = 1046574445
zrng(37) = 68911991
zrng(38) = 2088367019
zrng(39) = 748545416
zrng(40) = 622401386
zrng(41) = 2122378830
zrng(42) = 640690903
zrng(43) = 1774806513
zrng(44) = 2132545692
zrng(45) = 2079249579
zrng(46) = 78130110
zrng(47) = 852776735
zrng(48) = 1187867272
zrng(49) = 1351423507
zrng(50) = 1645973084
zrng(51) = 1997049139
zrng(52) = 922510944
zrng(53) = 2045512870
zrng(54) = 898585771
zrng(55) = 243649545
zrng(56) = 1004818771
zrng(57) = 773686062
zrng(58) = 403188473
zrng(59) = 372279877
zrng(60) = 1901633463
zrng(61) = 498067494
zrng(62) = 2087759558
zrng(63) = 493157915
zrng(64) = 597104727
zrng(65) = 1530940798
zrng(66) = 1814496276
zrng(67) = 536444882

```

zrng(68) = 1663153658
zrng(69) = 855503735
zrng(70) = 67784357
zrng(71) = 1432404475
zrng(72) = 619691088
zrng(73) = 119025595
zrng(74) = 880802310
zrng(75) = 176192644
zrng(76) = 1116780070
zrng(77) = 277854671
zrng(78) = 1366580350
zrng(79) = 1142483975
zrng(80) = 2026948561
zrng(81) = 1053920743
zrng(82) = 786262391
zrng(83) = 1792203830
zrng(84) = 1494667770
zrng(85) = 1923011392
zrng(86) = 1433700034
zrng(87) = 1244184613
zrng(88) = 1147297105
zrng(89) = 539712780
zrng(90) = 1545929719
zrng(91) = 190641742
zrng(92) = 1645390429
zrng(93) = 264907697
zrng(94) = 620389253
zrng(95) = 1502074852
zrng(96) = 927711160
zrng(97) = 364849192
zrng(98) = 2049576050
zrng(99) = 638580085
zrng(100) = 547070247

```

End Sub

Public Function lcggrand(Stream As Integer) As Double

```

Dim zi As Long
Dim lowprd As Long
Dim hi31 As Long

zi = zrng(Stream)

lowprd = (zi And 65535) * MULT1
hi31 = (zi \ 65536) * MULT1 + lowprd \ 65536
zi = ((lowprd And 65535) - MODLUS) + ((hi31 And 32767) * 65536) + (hi31 \ 32768)

If zi < 0 Then zi = zi + MODLUS

lowprd = (zi And 65535) * MULT2
hi31 = (zi \ 65536) * MULT2 + (lowprd \ 65536)

```

```

    zi = ((lowprd And 65535) - MODLUS) + ((hi31 And 32767) * 65536) + (hi31 \ 32768)

    If zi < 0 Then zi = zi + MODLUS

    zrng(Stream) = zi

    lcgrand = (zi \ 128 Or 1) / 16777216#

End Function

Sub lcgrandst(zset As Long, Stream As Integer)
    zrng(Stream) = zset
End Sub

Function lcgrandgt(Stream As Integer) As Long
    lcgrandgt = zrng(Stream)
End Function

```

A.3 Random-Variate Generation

```

Public Function Expon(Mean As Double, Stream As Integer) As Double
' function to generate exponential variates with mean Mean via inverse cdf

    Expon = -VBA.Log(1 - lcgrand(Stream)) * Mean

End Function

Public Function Uniform(Lower As Double, Upper As Double, _
                        Stream As Integer) As Double
' function to generate Uniform(Lower, Upper) variates via inverse cdf

    Uniform = Lower + (Upper - Lower) * lcgrand(Stream)

End Function

Public Function Random_integer(ByRef prob_distrib() As Double, _
                              Stream As Integer)
' function to generate a random integer

    Dim i As Integer
    Dim U As Double

    'Generate a U(0,1) random variate.
    U = lcgrand(Stream)

    'Return a random integer in accordance with the (cumulative) distribution
    ' function prob_distrib.
    Random_integer = 1
    While U >= prob_distrib(Random_integer)
        Random_integer = Random_integer + 1
    Wend

End Function

```

```

Public Function Erlang(m As Integer, Mean As Double, _
                    Stream As Integer) As Double
' function to generate an Erlang variate with given Mean and m phases

    Dim i As Integer
    Dim mean_exponential As Double
    Dim Sum As Double

    mean_exponential = Mean / m
    Sum = 0#
    For i = 1 To m
        Sum = Sum + Expon(mean_exponential, Stream)
    Next i
    Erlang = Sum
End Function

Public Function Triangular(a As Double, b As Double, c As Double, _
                    Stream As Integer) As Double
' function to generate triangular(a, b, c) variates via inverse
' note: differs from Law and Kelton in that a = min, b = mode, c = max

    Dim Standardb As Double
    Dim U As Double

    Standardb = (b - a) / (c - a)
    U = lcgrand(Stream)
    If U <= Standardb Then
        Triangular = VBA.Sqr(Standardb * U)
    Else
        Triangular = 1 - VBA.Sqr((1 - Standardb) * (1 - U))
    End If

    Triangular = a + (c - a) * Triangular

End Function

Public Function Normal(Mean As Double, Variance As Double, _
                    Stream As Integer) As Double
' function to generate normal variates via polar method

    Dim U1 As Double, U2 As Double, V1 As Double, V2 As Double, _
        W As Double, Y As Double

    Do
        U1 = lcgrand(Stream)
        U2 = lcgrand(Stream)
        V1 = 2 * U1 - 1
        V2 = 2 * U2 - 1
        W = V1 ^ 2 + V2 ^ 2
    Loop Until W <= 1

    Y = VBA.Sqr(-2 * VBA.Log(W) / W)
    Normal = V1 * Y
    Normal = Mean + VBA.Sqr(Variance) * Normal

End Function

```

```

Public Function Lognormal(MeanPrime As Double, VariancePrime As Double, _
                        Stream As Integer) As Double
' function to generate lognormal variate by transforming a normal
' note MeanPrime and VariancePrime as the DESIRED mean and variance
' for the lognormal

    Dim Mean As Double, Variance As Double

    Mean = VBA.Log(MeanPrime ^ 2 / VBA.Sqr(MeanPrime ^ 2 + VariancePrime))
    Variance = VBA.Log(1 + VariancePrime / MeanPrime ^ 2)

    Lognormal = VBA.Exp(Normal(Mean, Variance, Stream))

End Function

```

A.4 Class Modules

A.4.1 CTStat

```

' Generic continuous-time statistics object
' Note that CTStat should be called AFTER the value of the variable changes
' Last update 6/28/2010
Private Area As Double
Private Tlast As Double
Private Xlast As Double
Private TClear As Double

Private Sub Class_Initialize()
' Executes when CTStat object is created to initialize variables
    Area = 0
    Tlast = 0
    TClear = 0
    Xlast = 0
End Sub

Public Sub Record(X As Double)
' Update the CTStat from last time change and keep track of previous value
    Area = Area + Xlast * (Clock - Tlast)
    Tlast = Clock
    Xlast = X
End Sub

Function Mean() As Double
' Return the sample mean up through current time
' but do not update
    Mean = 0
    If (Clock - TClear) > 0 Then
        Mean = (Area + Xlast * (Clock - Tlast)) /
            (Clock - TClear)
    End If
End Function

Public Sub Clear()
' Clear statistics

```

```

    Area = 0
    Tlast = Clock
    TClear = Clock
End Sub

```

A.4.2 DTStat

```

' Generic discrete-time statistics object
' Last update 7/26/2009
Private Sum As Double
Private SumSquared As Double
Private NumberOfObservations As Double

Private Sub Class_Initialize()
' Executes when DTStat object is created to initialize variables
    Sum = 0
    SumSquared = 0
    NumberOfObseverations = 0
End Sub

Public Sub Record(X As Double)
' Update the DTStat
    Sum = Sum + X
    SumSquared = SumSquared + X * X
    NumberOfObservations = NumberOfObservations + 1
End Sub

Function Mean() As Double
' Return the sample mean
    Mean = 0
    If NumberOfObservations > 0 Then
        Mean = Sum / NumberOfObservations
    End If
End Function

Function StdDev() As Double
' Return the sample standard deviation
    StdDev = 0
    If NumberOfObservations > 1 Then
        StdDev = VBA.Sqrt((SumSquared - Sum ^ 2 _
            / NumberOfObservations) / (NumberOfObservations - 1))
    End If
End Function

Function N() As Double
' Return the number of observations collected
    N = NumberOfObservations
End Function

Public Sub Clear()
' Clear statistics
    Sum = 0
    SumSquared = 0
    NumberOfObservations = 0
End Sub

```

A.4.3 Entity

```
' This is a generic Entity that has a single attribute CreateTime
' Last update 6/28/2010
Public CreateTime As Double
' Add additional problem specific attributes here

Private Sub Class_Initialize()
' Executes when Entity object is created to initialize variables
    CreateTime = Clock
End Sub
```

A.4.4 EventCalendar

```
' This class module creates an Event Calendar object
' which is a list of event notices ordered by event time.
' Based on an object created by Steve Roberts @ NCSU
' Last update 7/26/2009
Private ThisCalendar As New Collection

Public Sub Schedule(addedEvent As EventNotice)
' Add EventNotice in EventTime order
    Dim i As Integer
    If ThisCalendar.Count = 0 Then 'no events in calendar
        ThisCalendar.Add addedEvent
    ElseIf ThisCalendar(ThisCalendar.Count).EventTime <= addedEvent.EventTime Then
        'added event after last event in calendar
        ThisCalendar.Add addedEvent, After:=ThisCalendar.Count
    Else
        'search for the correct place to insert the event
        For i = 1 To ThisCalendar.Count
            If ThisCalendar(i).EventTime > addedEvent.EventTime Then
                Exit For
            End If
        Next i
        ThisCalendar.Add addedEvent, before:=i
    End If
End Sub

Public Function Remove() As EventNotice
' Remove next event and return the EventNotice object
    If ThisCalendar.Count > 0 Then
        Set Remove = ThisCalendar.Item(1)
        ThisCalendar.Remove (1)
    End If
End Function

Function N() As Integer
' Return current number of events on the event calendar
    N = ThisCalendar.Count
End Function
```

A.4.5 EventNotice

```
' This is a generic EventNotice object with EventTime, EventType
' and WhichObject attributes
```

```

' Last update 7/26/2009
Public EventTime As Double
Public EventType As String
Public WhichObject As Object

' Add additional problem specific attributes here

```

A.4.6 FIFOQueue

```

' This is a generic FIFO Queue object that
' also keeps track of statistics on the number in queue (WIP)
' Last update 6/28/2010
Private ThisQueue As New Collection
Dim WIP As New CTStat

Private Sub Class_Initialize()
' Executes when FIFOQueue object is created to add queue statistics
' to TheCTStats collection
    TheCTStats.Add WIP
End Sub

Function NumQueue() As Integer
' Return current number in queue
    NumQueue = ThisQueue.Count
End Function

Public Sub Add(ThisEntity As Object)
' Add an entity to the end of the queue
    If ThisQueue.Count = 0 Then
        ThisQueue.Add ThisEntity
    Else
        ThisQueue.Add ThisEntity, After:=ThisQueue.Count
    End If
    WIP.Record ThisQueue.Count
End Sub

Public Function Remove() As Object
' Remove the first entity from the queue and return the object
' after updating queue statistics
    If ThisQueue.Count > 0 Then
        Set Remove = ThisQueue.Item(1)
        ThisQueue.Remove 1
    End If
    WIP.Record ThisQueue.Count
End Function

Public Function Mean() As Double
'Return the average number in queue up to the current time
    Mean = WIP.Mean
End Function

```

A.4.7 Resource

```

' This is a generic Resource object that
' also keeps track of statistics on the number of busy resources
' Last update 6/28/2010
Public Busy As Integer
Public NumberOfUnits As Integer
Dim NumBusy As New CTStat

Private Sub Class_Initialize()
' Executes when resource object is created to initialize variables
' and add number of busy Resources statistic to TheCTStats collection
    Busy = 0
    NumberOfUnits = 0
    TheCTStats.Add NumBusy
End Sub

Public Function Seize(Units As Integer) As Boolean
' Seize Units of resource then updates statistics
' Returns False and does not seize if not enough resources available;
' otherwise returns True
    Dim diff As Integer
    diff = NumberOfUnits - Units - Busy
' If diff is nonnegative, then there are enough resources to seize
    If diff >= 0 Then
        Busy = Busy + Units
        NumBusy.Record CDb1(Busy)
        Seize = True
    Else
        Seize = False
    End If
End Function

Public Function Free(Units As Integer) As Boolean
' Free Units of resource then updates statistics
' Returns False and does not free if attempting to free more
' resources than available; otherwise returns True
' CORRECTED 2/1/2010
    Dim diff As Integer
    diff = Busy - Units
' If diff is negative, then trying to free too many resources
    If diff < 0 Then
        Free = False
    Else
        Busy = Busy - Units
        NumBusy.Record CDb1(Busy)
        Free = True
    End If
End Function

Public Function Mean()
' Return time-average number of busy resources up to current time
    Mean = NumBusy.Mean
End Function

Public Sub SetUnits(Units As Integer)
' Set the capacity of the resource (number of identical units)
    NumberOfUnits = Units
End Sub

```

References

- Albright, S. C. (2007). *VBA for modelers: Developing decision support systems using Microsoft Excel* (2nd ed.). Belmont: Thompson Higher Education.
- Alexopoulos, C. (2006). Statistical estimation in computer simulation. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Andradóttir, S. (1999). Accelerating the convergence of random search methods for discrete stochastic optimization. *ACM Transactions on Modeling and Computer Simulation*, *9*, 349–380.
- Andradóttir, S. (2006a). An overview of simulation optimization via random search. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Andradóttir, S. (2006b). Simulation optimization with countably infinite feasible regions: Efficiency and convergence. *ACM Transactions on Modeling and Computer Simulation*, *16*, 357–374.
- Andradóttir, S., & Kim, S. (2010). Fully sequential procedures for comparing constrained systems via simulation. *Naval Research Logistics*, *57*, 403–421.
- Ankenman, B. E., & Nelson, B. L. (2012, in press). A quick assessment of input uncertainty. *Proceedings of the 2012 Winter Simulation Conference, Berlin*.
- Asmussen, S. & Glynn, P. W., (2007). *Stochastic simulation: Algorithms and analysis*. New York: Springer.
- Batur, D., & Kim, S. (2010). Finding feasible systems in the presence of constraints on multiple performance measures. *ACM Transactions on Modeling and Computer Simulation*, *20*, 13:1–13:26.
- Bechhofer, R. E., Santner, T. J., & Goldsman, D. (1995). *Design and analysis of experiments for statistical selection, screening and multiple comparisons*. New York: Wiley.
- Billar, B., & Corlu, C. G. (2012). Copula-based multivariate input modeling. *Surveys in Operations Research and Management Science*, *17*, 69–84.
- Billar, B., & Ghosh, S. (2006). Multivariate input processes. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Billar, B., & Nelson, B. L. (2003). Modeling and generating multivariate time-series input processes using a vector autoregressive technique. *ACM Transactions on Modeling and Computer Simulation*, *13*, 211–237.
- Billar, B., & Nelson, B. L. (2005). Fitting time series input processes for simulation. *Operations Research*, *53*, 549–559.
- Billingsley, P. (1995). *Probability and measure* (3rd ed.). New York: Wiley.

- Boesel, J., Nelson, B. L., & Kim, S. (2003). Using ranking and selection to “clean up” after simulation optimization. *Operations Research*, *51*, 814–825.
- Bratley, P., Fox, B. L., & Schrage, L. E. (1987). *A guide to simulation* (2nd ed.). New York: Springer.
- Burt, J. M., & Garman, M. B. (1971). Conditional Monte Carlo: A simulation technique for stochastic network analysis. *Management Science*, *19*, 207–217.
- Cario, M. C. & Nelson, B. L. (1998). Numerical methods for fitting and simulating autoregressive-to-anything processes. *INFORMS Journal on Computing*, *10*, 72–81.
- Cash, C., Nelson, B. L., Long, J., Dippold, D., & Pollard, W. (1992). Evaluation of tests for initial-condition bias. *Proceedings of the 1992 Winter Simulation Conference* (pp. 577–585). Piscataway, New Jersey: IEEE.
- Chatfield, C. (2004). *The analysis of time series: An introduction* (6th ed.). Boca Raton: Chapman & Hall/CRC.
- Chen, H. (2001). Initialization for NORTA: Generation of random vectors with specified marginals and correlations. *INFORMS Journal on Computing*, *13*, 312–331.
- Chow, Y. S., & Robbins, H. (1965). On the asymptotic theory of fixed-width sequential confidence intervals for the mean. *The Annals of Mathematical Statistics*, *36*, 457–462.
- Devroye, L. (1986). *Non-uniform random variate generation*. New York: Springer.
- Devroye, L. (2006). Nonuniform random variate generation. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. Boca Raton: Chapman & Hall/CRC.
- Elizandro, D., & Taha, H. (2008). *Simulation of industrial systems: Discrete event simulation using Excel/VBA*. New York: Auerbach Publications.
- Frazier, P. I. (2010). Decision-theoretic foundations of simulation optimization. In J. J. Cochran (Ed.), *Wiley encyclopedia of operations research and management sciences*. New York: Wiley.
- Fu, M. C. (2006). Gradient estimation. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Gerhardt, I., & Nelson, B. L. (2009). Transforming renewal processes for simulation of nonstationary arrival processes. *INFORMS Journal on Computing*, *21*, 630–640.
- Ghosh, S., & Henderson, S. G. (2002). Chessboard distributions and random vectors with specified marginals and covariance matrix. *Operations Research*, *50*, 820–834.
- Glasserman, P. (2004). *Monte Carlo methods in financial engineering*. New York: Springer.
- Glasserman, P., & Yao, D. D. (1992). Some guidelines and guarantees for common random numbers. *Management Science*, *38*, 884–908.
- Glynn, P. W. (2006). Simulation algorithms for regenerative processes. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Glynn, P. W., & Whitt, W. (1992). The asymptotic validity of sequential stopping rules for stochastic simulations. *The Annals of Applied Probability*, *2*, 180–198.
- Goldman, D., Kim, S., Marshall, S. W., & Nelson, B. L. (2002). Ranking and selection for steady-state simulation: Procedures and perspectives. *INFORMS Journal on Computing*, *14*, 2–19.
- Goldman, D., & Nelson, B. L. (1998). Comparing systems via simulation. In J. Banks (Ed.), *Handbook of simulation* (pp. 273–306). New York: Wiley.
- Goldman, D., & Nelson, B. L. (2006). Correlation-based methods for output analysis. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Gross, D., Shortle, J. F., Thompson, J. M., & Harris, C. M. (2008). *Fundamentals of queueing theory* (4th ed.). New York: Wiley.
- Haas, P. J. (2002). *Stochastic petri nets: Modeling, stability, simulation*. New York: Springer.
- Henderson, S. G. (2003). Estimation of nonhomogeneous Poisson processes from aggregated data. *Operations Research Letters*, *31*, 375–382.

- Henderson, S. G. (2006). Mathematics for simulation. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Henderson, S. G., & Nelson, B. L. (2006). Stochastic computer simulation. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Hill, R. R., & Reilly, C. H. (2000). The effects of coefficient correlation structure in two-dimensional knapsack problems on solution procedure performance. *Management Science*, *46*, 302–317.
- Hong, L. J., & Nelson, B. L. (2007a). Selecting the best system when systems are revealed sequentially. *IIE Transactions*, *39*, 723–734.
- Hong, L. J., & Nelson, B. L. (2007b). A framework for locally convergent random-search algorithms for discrete optimization via simulation. *ACM Transactions on Modeling and Computer Simulation*, *17*, 19/1-19/22.
- Hörmann, W. (1993). The transformed rejection method for generating Poisson random variables. *Insurance: Mathematics and Economics*, *12*, 39–45.
- Iravani, S. M. R., & Krishnamurthy, V. (2007). Workforce agility in repair and maintenance environments. *Manufacturing and Service Operations Management*, *9*, 168–184.
- Iravani, S. M., Van Oyen, M. P., & Sims, K. T. (2005). Structural flexibility: A new perspective on the design of manufacturing and service operations. *Management Science*, *51*, 151–166.
- Johnson, M. E. (1987). *Multivariate statistical simulation*. New York: Wiley.
- Johnson, N. L., Kemp, A. W., & Kotz, S. (2005). *Univariate discrete distributions* (3rd ed.). New York: Wiley.
- Johnson, N. L., Kotz, S., & Balakrishnan, N. (1994). *Continuous univariate distributions* (2nd ed., Vol. 1). New York: Wiley.
- Johnson, N. L., Kotz, S., & Balakrishnan, N. (1995). *Continuous univariate distributions* (2nd ed., Vol. 2). New York: Wiley.
- Johnson, N. L., Kotz, S., & Balakrishnan, N. (1997). *Discrete multivariate distributions*. New York: Wiley.
- Kachitvichyanukul, V., & Schmeiser, B. (1990). Noninverse correlation induction: Guidelines for algorithm development. *Journal of Computational and Applied Mathematics*, *31*, 173–180.
- Karian, A. Z., & Dudewicz, E. J. (2000). *Fitting statistical distributions: The generalized lambda distribution and generalized bootstrap methods*. New York: CRC.
- Karlin, S., & Taylor, H. M. (1975). *A first course in stochastic processes* (2nd ed.). New York: Academic.
- Kelton, W. D. (2006). Implementing representations of uncertainty. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Kelton, W. D., Smith, J. S., & Sturrock, D. T. (2011). *Simio and simulation: Modeling, analysis and applications* (2nd ed.). New York: McGraw-Hill.
- Kim, S., & Nelson, B. L. (2001). A fully sequential procedure for indifference-zone selection in simulation. *ACM Transactions on Modeling and Computer Simulation*, *11*, 251–273.
- Kim, S., & Nelson, B. L. (2006). Selecting the best system. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Knuth, D. E. (1998). *The art of computer programming, Vol. 2: Seminumerical algorithms* (3rd ed.). Boston: Addison-Wesley.
- Kotz, S., Balakrishnan, N., & Johnson, N. L. (2000). *Continuous multivariate distributions, Vol. 1, models and applications* (2nd ed.). New York: Wiley.
- Kulkarni, V. G. (1995). *Modeling and analysis of stochastic systems*. London: Chapman & Hall.
- Lakhany, A., & Mausser, H. (2000). Estimating the parameters of the generalized lambda distribution. *ALGO Research Quarterly*, *3*, 47–58.
- Law, A. M. (2007). *Simulation modeling and analysis* (4th ed.). New York: McGraw-Hill.

- Law, A. M., & Kelton, W. D. (2000). *Simulation modeling and analysis* (3rd ed.). New York: McGraw-Hill.
- L'Ecuyer, P. (1988). Efficient and portable combined random number generators. *Communications of the ACM*, *31*, 742–749.
- L'Ecuyer, P. (1990). A unified view of IPA, SF, and LR gradient estimation techniques. *Management Science*, *36*, 1364–1383.
- L'Ecuyer, P. (1999). Good parameters and implementations for combined multiple recursive random number generators. *Operations Research*, *47*, 159–164.
- L'Ecuyer, P. (2006). Uniform random number generation. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- L'Ecuyer, P., & Simard, R. (2001). On the performance of birthday spacings tests for certain families of random number generators. *Mathematics and Computers in Simulation*, *55*, 131–137.
- L'Ecuyer, P., Simard, R., Chen, E. J., & Kelton, W. D. (2002). An object-oriented random-number package with many long streams and substreams. *Operations Research*, *50*, 1073–1075.
- Lee, S., Wilson, J. R., & Crawford, M. M. (1991). Modeling and simulation of a nonhomogeneous Poisson process having cyclic behavior. *Communications in Statistics-Simulation and Computation*, *20*, 777–809.
- Leemis, L. M. (1991). Nonparametric estimation of the cumulative intensity function for a non-homogeneous Poisson process. *Management Science*, *37*, 886–900.
- Leemis, L. M. (2006). Arrival processes, random lifetimes and random objects. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Leemis, L. M., & McQueston, J. T. (2008). Univariate distribution relationships. *The American Statistician*, *62*, 45–53.
- Lehmann, E. L. (2010). *Elements of large-sample theory*. New York: Springer.
- Lewis, T. A. (1981). Confidence intervals for a binomial parameter after observing no successes. *The American Statistician*, *35*, 154.
- Marse, K., & Roberts, S. D. (1983). Implementing a portable FORTRAN uniform (0,1) generator. *Simulation*, *41*, 135–139.
- Montgomery, D. C. (2009). *Design and analysis of experiments* (7th ed.). New York: Wiley.
- Nádas, A. (1969). An extension of a theorem of Chow and Robbins on sequential confidence intervals for the mean. *The Annals of Mathematical Statistics*, *40*, 667–671.
- Nelson, B. L. (1990). Control-variate remedies. *Operations Research*, *38*, 974–992.
- Nelson, B. L. (1995). *Stochastic modeling: Analysis and simulation*. Mineola: Dover Publications, Inc.
- Nelson, B. L. (2008). The MORE plot: Displaying measures of risk and error from simulation output. *Proceedings of the 2008 Winter Simulation Conference* (pp. 413–416). Piscataway, New Jersey: IEEE.
- Nelson, B. L., & Taaffe, M. R. (2004). The $Ph_t/Ph_t/\infty$ queueing system: Part I: The single node. *INFORMS Journal on Computing*, *16*, 266–274.
- Nelson, B. L., Swann, J., Goldsman, D., & Song, W. (2001). Simple procedures for selecting the best simulated system when the number of alternatives is large. *Operations Research*, *49*, 950–963.
- Pasupathy, R., & Schmeiser, B. (2010). The initial transient in steady-state point estimation: Contexts, a bibliography, the MSE criterion, and the MSER statistic. *Proceedings of the 2010 Winter Simulation Conference* (pp. 184–197). Piscataway, New Jersey: IEEE.
- Sargent, R. G. (2011). Verification and validation of simulation models. *Proceedings of the 2011 Winter Simulation Conference* (pp. 183–198). Piscataway, New Jersey: IEEE.
- Schmeiser, B. (1982). Batch size effects in the analysis of simulation output. *Operations Research*, *30*, 556–568.
- Schruben, L. (1982). Detecting initialization bias in simulation output. *Operations Research*, *30*, 569–590.

- Schruben, L., Singh, H., & Tierney, L. (1983). Optimal tests for initialization bias in simulation output. *Operations Research*, *31*, 1167–1178.
- Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2009). *Lectures on stochastic programming: Modeling and theory*. Philadelphia: Society for Industrial and Applied Mathematics.
- Shechter, S. M., Schaefer, A. J., Braithwaite, R. S., & Roberts, M. S. (2006). Increasing the efficiency of Monte Carlo cohort simulations with variance reduction techniques. *Medical Decision Making*, *26*, 550–553.
- Snell, M., & Schruben, L. (1985). Weighting simulation data to reduce initialization effects. *IIE Transactions*, *17*, 354–363.
- Steiger, N. M., & Wilson, J. R. (2001). Convergence properties of the batch means method for simulation output analysis. *INFORMS Journal on Computing*, *13*, 277–293.
- Stigler, S. M. (1986). *The history of statistics: The measurement of uncertainty before 1900*. Cambridge, MA: Belknap.
- Swain, J. J., Venkatraman, S., & Wilson, J. R. (1988). Least-squares estimation of distribution functions in Johnson's translation system. *Journal of Statistical Computation and Simulation*, *29*, 271–297.
- Tafazzoli, A., & Wilson, J. R. (2011). Skart: A skewness-and-autoregression-adjusted batch-means procedure for simulation analysis. *IIE Transactions*, *43*, 110–128.
- Walkenbach, J. (2010). *Excel 2010 power programming with VBA*. New York: Wiley.
- White, K. P. (1997). An effective truncation heuristic for bias reduction in simulation output. *Simulation*, *69*, 323–334.
- Whitt, W. (1981). Approximating a point process by a renewal process: The view through a queue, an indirect approach. *Management Science*, *27*, 619–636.
- Whitt, W. (1989). Planning queueing simulations. *Management Science*, *35*, 1341–1366.
- Whitt, W. (2006). Analysis for design. In S. G. Henderson & B. L. Nelson (Eds.), *Handbooks in operations research and management science: Simulation*. New York: North-Holland.
- Whitt, W. (2007). What you should know about queueing models to set staffing requirements in service systems. *Naval Research Logistics*, *54*, 476–484.
- Xu, J., Hong, L. J., & Nelson, B. L. (2010). Industrial strength COMPASS: A comprehensive algorithm and software for optimization via simulation. *ACM Transactions on Modeling and Computer Simulation*, *20*, 1–29.
- Xu, J., Nelson, B. L., & Hong, L. J. (2012, in press). An adaptive hyperbox algorithm for high-dimensional discrete optimization via simulation problems. *INFORMS Journal on Computing*.

Index

- $M(t)/M/\infty$ queue, 29
- $M/G/1$ queue, 32, 49
- $M/M/\infty$, 163, 167
- p -value, 111
- t distribution, 51

- absolute error, 178, 179
- across-replication output, 4
- adaptive random search, 214
- Albright, S. C., 7, 267
- Alexopoulos, C., 179, 267
- Andradóttir, S., 197, 214, 227, 267
- Ankenman, B. E., 166, 267
- arrival rate, 124
- arrival-counting process, 106, 123
- Asian option, 38, 64
- Asmussen, S., viii, 230, 267
- asymptotic behavior, 83
- asymptotic bias, 91
- asymptotic MSE, 183
- asymptotic variance, 33, 91
- autocorrelation, 91

- Balakrishnan, N., 105, 231, 269
- batch means, 190, 191
- batch size, 190
- batching, 181, 192
- Batur, D., 227, 267
- Bechhofer, R. E., 206, 267
- beta distribution, 137
- Biller, B., 122, 123, 133, 146–148, 267
- Billingsley, P., 90, 267
- binomial distribution, 105
- Black–Scholes formula, 38
- Boesel, J., 200, 208, 213, 268
- Bonferroni inequality, 199
- bootstrapping, 165

- Braithwaite, R. S., 204, 271
- Bratley, P., 118, 149, 230, 268
- Burt, J. M., 35, 268

- Cario, M. C., 148, 268
- Cash, C., 187, 268
- central difference, 221
- central limit theorem, 87, 88, 105, 160
- central moment, 113
- Chatfield, C., 34, 268
- Chen, E. J., 204, 270
- Chen, H., 148, 268
- Cholesky decomposition, 147
- Chow, Y. S., 179, 268
- clock, 2
- coefficient of variation, 249
- common random numbers, 200, 201, 230, 245
- confidence interval, 87, 93, 159, 178
- continuous mapping theorem, 93
- control variates, 230, 235, 251
- convergence in distribution, 33, 87
- convergence in mean square, 97
- convergence in probability, 86
- convergence with probability 1, 86
- converging together lemma, 92
- copula, 133
- Corlu, C. G., 133, 267
- correlation-matching problem, 147
- Crawford, M. M., 150, 270

- Delta method, 93
- Dentcheva, D., 224, 271
- Devroye, L., 135, 268
- Dippold, D., 187, 268
- distribution bounds, 112
- doubly stochastic Poisson process, 152
- Dudewicz, E. J., 116, 144, 269

- Efron, B., 165, 268
 Elizandro, D., 7, 268
 empirical cdf (ecdf), 47, 117
 empirical distribution, 103
 equilibrium renewal process, 124
 Erlang distribution, 49, 167
 estimation error, 83
 event, 2
 event calendar, 2
 exact random-variate generation, 14
 experiment, 80, 244
 experiment design, 175, 176
 exponential distribution, 12, 29

 finite difference, 217
 Fox, B. L., 118, 149, 230, 268
 Frazier, P. I., 206, 212, 268
 Fu, M. C., 217, 221, 268

 gamma distribution, 108
 Garman, M. B., 35, 268
 generalized lambda distribution (GLD), 116, 144
 geometric Brownian motion, 38, 64
 Gerhardt, I., 126, 132, 153, 268
 Ghosh, S., 122, 123, 146, 267, 268
 Glasserman, P., 37, 158, 203, 218, 230, 234, 268
 global convergence, 197
 Glynn, P. W., viii, 90, 94, 179, 230, 267, 268
 Goldsman, D., 191, 206, 208, 244, 247, 267, 268, 270
 goodness-of-fit test, 110
 gradient estimation, 217
 Gross, D., 32, 33, 268

 Hörmann, W., 138, 269
 Haas, P. J., 90, 94, 268
 Harris, C. M., 32, 33, 268
 hazard function, 108
 Henderson, S. G., 35, 90, 130, 146, 195, 268, 269
 Hill, R. R., 245, 269
 Hong, L. J., 213–215, 269, 271
 hyperexponential distribution, 153

 i.i.d., 4
 indifference-zone selection, 209
 infinitesimal perturbation analysis, 218
 input, 2
 input modeling, 82, 99
 input uncertainty, 83, 163, 164, 167
 inputs, 80
 integrated rate function, 124

 interarrival time, 123
 inverse cdf, 13, 46
 inverse gamma distribution, 167
 Iravani, S. M. R., vii, 246, 269

 Java, viii, 16
 Johnson, M. E., 146, 269
 Johnson, N. L., 105, 231, 269

 Kachitvichyanukul, V., 205, 269
 Karian, A. Z., 116, 144, 269
 Karlin, S., 107, 177, 269
 Kelton, W. D., 171, 204, 269, 270
 Kemp, A. W., 105, 269
 Kim, S., 200, 206, 208, 209, 213, 227, 244, 267–269
 knapsack problem, 244
 Knuth, D. E., 179, 269
 Kolmogorov–Smirnov test, 105
 Kotz, S., 105, 231, 269
 Krishnamurthy, V., vii, 269
 Kulkarni, V. G., 124, 269
 kurtosis, 113

 L'Ecuyer, P., 139, 140, 142, 204, 217, 221, 270
 Lakhany, A., 116, 144, 269
 large-sample behavior, 83
 Law, A. M., viii, 41, 105, 109, 139, 230, 269, 270
 lazy simulator's rule, 2
 least-squares fit, 115
 Lee, S., 150, 270
 Leemis, L. M., 122, 128, 270
 Lehmann, E. L., 105, 270
 Lewis, T. A., 160, 270
 likelihood ratio, 220
 Lindley's equation, 32, 49, 90
 linearly interpolated ecdf, 118
 local convergence, 198
 locally optimal, 198
 logic, 80
 lognormal distribution, 107
 Long, J., 187, 268

 macroreplication, 251
 marginal standard error rule (MSER), 188, 237
 Marse, K., 139, 270
 Marshall, W. S., 244, 268
 Matlab, viii, 16
 Mausser, H., 116, 144, 269
 maximum likelihood estimation (MLE), 99, 100, 104, 110
 McQueston, J. T., 270
 mean squared error, 92

- measure of risk and error (MORE) plot, 161
 method of moments, 114
 mixture of Erlangs, 153
 modeling error, 82
 moment, 113
 Montgomery, D. C., 165, 246, 270
 multiple recursive generator, 141
 multiplicative congruential generator, 139
 multivariate input, 103, 122
 multivariate normal random vectors, 146
- Nádas, A., 179, 270
 negative binomial distribution, 105
 Nelson, B. L., 30, 35, 126, 132, 146–148, 153, 162, 166, 187, 191, 195, 200, 206, 208, 209, 213–215, 232, 235, 244, 247, 251, 267–271
 nonstationary arrival process, 122–124
 nonstationary Poisson process, 29, 44, 66, 76, 100, 131
 normal distribution, 105
 normal to anything (NORTA), 133, 145
- option, 37
 order statistics, 115, 157
 output, 2
 output analysis, 82
- paired- t confidence interval, 205
 parametric distribution, 103
 Pasupathy, R., 188, 270
 penalty function, 228
 performance parameters, 80
 period, 14, 139
 Poisson distribution, 105
 Poisson process, 100
 Pollaczek–Khinchine formula, 33, 84, 102, 249
 Pollard, W., 187, 268
 practitioner’s experiment, 243
 probability–integral transformation, 153
 pseudorandom numbers, 14, 138, 201
 pseudorandom-number generator, 139, 204
- quantile, 157, 181
 quantile confidence interval, 161
 quantile–quantile (Q-Q) plot, 111
- random effects model, 165
 random numbers, 11, 14
 random test problem, 244
 random variates, 11, 134
 random vectors, 122, 133
 random-number generation, 99
 random-number stream, 42, 140
 random-variate generation, 99
 ranking and selection, 244
 regenerative process, 94
 Reilly, C. H., 245, 269
 rejection method, 135
 relative error, 39, 179
 renewal process, 94, 105, 123
 replication, 4, 11, 15
 researcher’s experiment, 243
 Robbins, H., 179, 268
 Roberts, M. S., 204, 271
 Roberts, S. D., 139, 270
 Ruczynski, A., 224, 271
- sample average approximation, 223
 sample path, 2
 Santner, T. J., 206, 267
 Sargent, R. G., 82, 270
 scenario variable, 175
 Schaefer, A. J., 204, 271
 Schmeiser, B., 188, 191, 205, 269, 270
 Schrage, L. E., 118, 149, 230, 268
 Schruben, L., 187, 236, 270, 271
 Shapiro, A., 224, 271
 Shechter, S. M., 204, 271
 Shortle, J. F., 32, 33, 268
 Simard, R., 140, 204, 270
 simlib, 41
 Sims, K. T., 246, 269
 simulation optimization, 194
 Singh, H., 187, 271
 skewness, 113
 Smith, J. S., 171, 269
 Snell, M., 236, 271
 Song, W., 208, 247, 270
 standard error, 33, 59, 92
 state, 2
 steady-state, fixed-budget problem, 182
 steady-state, fixed-precision problem, 182
 Steiger, N. M., 191, 271
 Stigler, S. M., 177, 271
 stochastic activity network (SAN), 35, 58
 stochastic approximation, 222
 stochastic constraints, 225
 stochastic process, 83
 stochastic simulation, 1
 stopping time, 175, 177
 streams, 143, 201, 204
 strong law of large numbers, 4, 88
 Sturrock, D. T., 171, 269
 superposition, 106
 Swain, J. J., 116, 271
 Swann, J., 208, 247, 270

- synchronization, 203
- system, 80

- Taaffe, M. R., 30, 270
- Tafazzoli, A., 191, 271
- Taha, H., 7, 268
- Taylor, H. M., 107, 177, 269
- thinning, 46, 76, 125
- Thompson, J. M., 32, 33, 268
- Tibshirani, R. J., 165, 268
- tie-breaking rule, 5, 6, 26
- Tierney, L., 187, 271
- time average, 4
- time series input process, 123
- traffic intensity, 249
- triangular distribution, 121
- truncated average, 49, 182

- uniform convergence, 224
- uniform distribution, 12
- univariate input, 103, 104

- validation, 82
- Van Oyen, M. P., 246, 269
- variance reduction technique, 230
- variate generation, 82
- VBA, viii, 7, 16
- VBASim, 16, 41, 253
- Venkatraman, S., 116, 271

- Walkenbach, J., 16, 271
- weak derivative, 220
- weak law of large numbers, 88
- Weibull distribution, 108
- White, K. P., 188, 271
- Whitt, W., 30, 34, 91, 179, 250, 268, 271
- Wilson, J. R., 116, 150, 191, 270, 271
- within-replication output, 4

- Xu, J., 213, 215, 271

- Yao, D. D., 203, 218, 268