华章教育

ELSEVIER 爱思唯尔

华章统计学原版精品系列

统计模拟

Simulation (Fifth Edition)

（英文版·第5版）

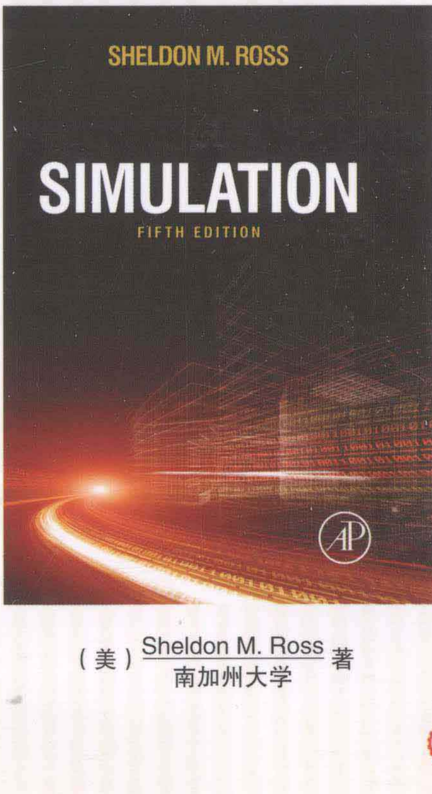

（美）Sheldon M. Ross 著
南加州大学

机械工业出版社
China Machine Press

华章统计学原版精品系列

统计模拟

Simulation (Fifth Edition)

（英文版·第5版）

（美）Sheldon M. Ross 著
南加州大学

机械工业出版社
China Machine Press

图书在版编目（CIP）数据

统计模拟（英文版·第5版）/（美）罗斯（Ross, S. M.）著. —北京：机械工业出版社，2013.5
（华章统计学原版精品系列）
书名原文：Simulation

ISBN 978-7-111-42045-3

I. 统… II. 罗… III. 统计－模拟实验－英文 IV. C8-33

中国版本图书馆CIP数据核字（2013）第066956号

本书版权登记号：图字：01-2013-1801

Sheldon M. Ross：Simulation, Fifth Edition (ISBN: 978-0-12-415825-2).

Elsevier (Singapore) Pte Ltd.
3 Killiney Road
#08-01 Winsland House I
Singapore 239519
Tel: (65) 6349-0200
Fax: (65) 6733-1817
First Published 2013

机械工业出版社（北京市西城区百万庄大街22号 邮政编码 100037）
责任编辑：明永玲
北京京师印务有限公司印刷
2013年5月第1版第1次印刷
147mm × 210mm · 10印张
标准书号：ISBN 978-7-111-42045-3
定 价：59.00元

凡购本书，如有缺页、倒页、脱页，由本社发行部调换
客服热线：(010) 88378991 88361066 投稿热线：(010) 88379604
购书热线：(010) 68326294 88379649 68995259 读者信箱：hzjsj@hzbook.com

Preface

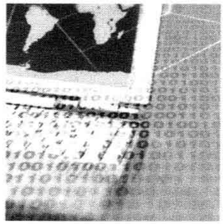

Overview

In formulating a stochastic model to describe a real phenomenon, it used to be that one compromised between choosing a model that is a realistic replica of the actual situation and choosing one whose mathematical analysis is tractable. That is, there did not seem to be any payoff in choosing a model that faithfully conformed to the phenomenon under study if it were not possible to mathematically analyze that model. Similar considerations have led to the concentration on asymptotic or steady-state results as opposed to the more useful ones on transient time. However, the advent of fast and inexpensive computational power has opened up another approach—namely, to try to model the phenomenon as faithfully as possible and then to rely on a simulation study to analyze it.

In this text we show how to analyze a model by use of a simulation study. In particular, we first show how a computer can be utilized to generate random (more precisely, pseudorandom) numbers, and then how these random numbers can be used to generate the values of random variables from arbitrary distributions. Using the concept of discrete events we show how to use random variables to generate the behavior of a stochastic model over time. By continually generating the behavior of the system we show how to obtain estimators of desired quantities of interest. The statistical questions of when to stop a simulation and what confidence to place in the resulting estimators are considered. A variety of ways in which one can improve on the usual simulation estimators are presented. In addition, we show how to use simulation to determine whether the stochastic model chosen is consistent with a set of actual data.

New to This Edition

- New exercises in most chapters.
- A new Chapter 6, dealing both with the multivariate normal distribution, and with copulas, which are useful for modeling the joint distribution of random variables.
- Chapter 9, dealing with variance reduction, includes additional material on stratification. For instance, it is shown that stratifying on a variable always results in an estimator having smaller variance than would be obtaind by using that variable as a control. There is also a new subsection on the use of post stratification.
- There is a new chapter dealing with additional variance reduction methods beyond those previously covered. Chapter 10 introduces the conditional Bernoulli sampling method, normalized importance sampling, and Latin Hypercube sampling.
- The chapter on Markov chain Monte Carlo methods has an new section entitled *Continuous time Markov chains and a Queueing Loss Model.*

Chapter Descriptions

The successive chapters in this text are as follows. **Chapter 1** is an introductory chapter which presents a typical phenomenon that is of interest to study. **Chapter 2** is a review of probability. Whereas this chapter is self-contained and does not assume the reader is familiar with probability, we imagine that it will indeed be a review for most readers. **Chapter 3** deals with random numbers and how a variant of them (the so-called pseudorandom numbers) can be generated on a computer. The use of random numbers to generate discrete and then continuous random variables is considered in Chapters 4 and 5.

Chapter 6 studies the multivariate normal distribution, and introduces copulas which are useful for modeling the joint distribution of random variables. **Chapter 7** presents the discrete event approach to track an arbitrary system as it evolves over time. A variety of examples—relating to both single and multiple server queueing systems, to an insurance risk model, to an inventory system, to a machine repair model, and to the exercising of a stock option—are presented. **Chapter 8** introduces the subject matter of statistics. Assuming that our average reader has not previously studied this subject, the chapter starts with very basic concepts and ends by introducing the bootstrap statistical method, which is quite useful in analyzing the results of a simulation.

Chapter 9 deals with the important subject of variance reduction. This is an attempt to improve on the usual simulation estimators by finding ones having the same mean and smaller variances. The chapter begins by introducing the technique of using antithetic variables. We note (with a proof deferred to the chapter's appendix) that this always results in a variance reduction along with a computational savings when we are trying to estimate the expected value of a function that is monotone in each of its variables. We then introduce control variables and illustrate their usefulness in variance reduction. For instance, we show how control variables can be effectively utilized in analyzing queueing systems, reliability systems, a list reordering problem, and blackjack. We also indicate how to use regression packages to facilitate the resulting computations when using control variables. Variance reduction by use of conditional expectations is then considered, and its use is indicated in examples dealing with estimating π, and in analyzing finite capacity queueing systems. Also, in conjunction with a control variate, conditional expectation is used to estimate the expected number of events of a renewal process by some fixed time. The use of stratified sampling as a variance reduction tool is indicated in examples dealing with queues with varying arrival rates and evaluating integrals. The relationship between the variance reduction techniques of conditional expectation and stratified sampling is explained and illustrated in the estimation of the expected return in video poker. Applications of stratified sampling to queueing systems having Poisson arrivals, to computation of multidimensional integrals, and to compound random vectors are also given. The technique of importance sampling is next considered. We indicate and explain how this can be an extremely powerful variance reduction technique when estimating small probabilities. In doing so, we introduce the concept of tilted distributions and show how they can be utilized in an importance sampling estimation of a small convolution tail probability. Applications of importance sampling to queueing, random walks, and random permutations, and to computing conditional expectations when one is conditioning on a rare event are presented. The final variance reduction technique of Chapter 9 relates to the use of a common stream of random numbers. **Chapter 10** introduces additional variance reduction techniques.

Chapter 11 is concerned with statistical validation techniques, which are statistical procedures that can be used to validate the stochastic model when some real data are available. Goodness of fit tests such as the chi-square test and the Kolmogorov–Smirnov test are presented. Other sections in this chapter deal with the two-sample and the n-sample problems and with ways of statistically testing the hypothesis that a given process is a Poisson process.

Chapter 12 is concerned with Markov chain Monte Carlo methods. These are techniques that have greatly expanded the use of simulation in recent years. The standard simulation paradigm for estimating $\theta = E[h(\mathbf{X})]$, where $\mathbf{X}$ is a random vector, is to simulate independent and identically distributed copies of $\mathbf{X}$ and then use the average value of $h(\mathbf{X})$ as the estimator. This is the so-called "raw" simulation estimator, which can then possibly be improved upon by using one or more of the variance reduction ideas of Chapters 9 and 10. However, in order to employ this approach it is necessary both that the distribution of $\mathbf{X}$ be specified and also that we be able to simulate from this distribution. Yet, as we see in Chapter 12, there are many examples where the distribution of $\mathbf{X}$ is known but we are not able to directly simulate the random vector $\mathbf{X}$, and other examples where the distribution is not completely known but is only specified up to a multiplicative constant. Thus, in either case, the usual approach to estimating θ is not available. However, a new approach, based on generating a Markov chain whose limiting distribution is the distribution of $\mathbf{X}$, and estimating θ by the average of the values of the function h evaluated at the successive states of this chain, has become widely used in recent years. These Markov chain Monte Carlo methods are explored in Chapter 12. We start, in Section 12.2, by introducing and presenting some of the properties of Markov chains. A general technique for generating a Markov chain having a limiting distribution that is specified up to a multiplicative constant, known as the Hastings–Metropolis algorithm, is presented in Section 12.3, and an application to generating a random element of a large "combinatorial" set is given. The most widely used version of the Hastings–Metropolis algorithm is known as the Gibbs sampler, and this is presented in Section 12.4. Examples are discussed relating to such problems as generating random points in a region subject to a constraint that no pair of points are within a fixed distance of each other, to analyzing product form queueing networks, to analyzing a hierarchical Bayesian statistical model for predicting the numbers of home runs that will be hit by certain baseball players, and to simulating a multinomial vector conditional on the event that all outcomes occur at least once. An application of the methods of this chapter to deterministic optimization problems, called simulated annealing, is presented in Section 12.5, and an example concerning the traveling salesman problem is presented. The final section of Chapter 12 deals with the sampling importance resampling algorithm, which is a generalization of the acceptance–rejection technique of Chapters 4 and 5. The use of this algorithm in Bayesian statistics is indicated.

Thanks

We are indebted to Yontha Ath (California State University, Long Beach) David Butler (Oregon State University), Matt Carlton (California Polytechnic State University), James Daniel (University of Texas, Austin), William Frye (Ball State University), Mark Glickman (Boston University), Chuanshu Ji (University of North Carolina), Yonghee Kim-Park (California State University, Long Beach), Donald E. Miller (St. Mary's College), Krzysztof Ostaszewski (Illinois State University), Bernardo Pagnocelli, Erol Peköz (Boston University), Yuval Peres (University of California, Berkeley), John Grego (University of South Carolina, Columbia), Zhong Guan (Indiana University, South Bend), Nan Lin (Washington University in St. Louis), Matt Wand (University of Technology, Sydney), Lianming Wang (University of South Carolina, Columbia), and Esther Portnoy (University of Illinois, Urbana-Champaign) for their many helpful comments. We would like to thank those text reviewers who wish to remain anonymous.

Contents

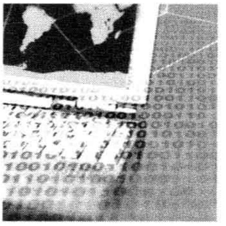

Introduction

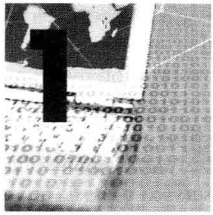

Consider the following situation faced by a pharmacist who is thinking of setting up a small pharmacy where he will fill prescriptions. He plans on opening up at 9 A.M. every weekday and expects that, on average, there will be about 32 prescriptions called in daily before 5 P.M. experience that the time that it will take him to fill a prescription, once he begins working on it, is a random quantity having a mean and standard deviation of 10 and 4 minutes, respectively. He plans on accepting no new prescriptions after 5 P.M., although he will remain in the shop past this time if necessary to fill all the prescriptions ordered that day. Given this scenario the pharmacist is probably, among other things, interested in the answers to the following questions:

1. What is the average time that he will depart his store at night?
2. What proportion of days will he still be working at 5:30 P.M.?
3. What is the average time it will take him to fill a prescription (taking into account that he cannot begin working on a newly arrived prescription until all earlier arriving ones have been filled)?
4. What proportion of prescriptions will be filled within 30 minutes?
5. If he changes his policy on accepting all prescriptions between 9 A.M. and 5 P.M., but rather only accepts new ones when there are fewer than five prescriptions still needing to be filled, how many prescriptions, on average, will be lost?
6. How would the conditions of limiting orders affect the answers to questions 1 through 4?

In order to employ mathematics to analyze this situation and answer the questions, we first construct a probability model. To do this it is necessary to

Simulation. DOI: http://dx.doi.org/10.1016/B978-0-12-415825-2.00001-2

make some reasonably accurate assumptions concerning the preceding scenario. For instance, we must make some assumptions about the probabilistic mechanism that describes the arrivals of the daily average of 32 customers. One possible assumption might be that the arrival rate is, in a probabilistic sense, constant over the day, whereas a second (probably more realistic) possible assumption is that the arrival rate depends on the time of day. We must then specify a probability distribution (having mean 10 and standard deviation 4) for the time it takes to service a prescription, and we must make assumptions about whether or not the service time of a given prescription always has this distribution or whether it changes as a function of other variables (e.g., the number of waiting prescriptions to be filled or the time of day). That is, we must make probabilistic assumptions about the daily arrival and service times. We must also decide if the probability law describing a given day changes as a function of the day of the week or whether it remains basically constant over time. After these assumptions, and possibly others, have been specified, a probability model of our scenario will have been constructed.

Once a probability model has been constructed, the answers to the questions can, in theory, be analytically determined. However, in practice, these questions are much too difficult to determine analytically, and so to answer them we usually have to perform a simulation study. Such a study programs the probabilistic mechanism on a computer, and by utilizing "random numbers" it simulates possible occurrences from this model over a large number of days and then utilizes the theory of statistics to estimate the answers to questions such as those given. In other words, the computer program utilizes random numbers to generate the values of random variables having the assumed probability distributions, which represent the arrival times and the service times of prescriptions. Using these values, it determines over many days the quantities of interest related to the questions. It then uses statistical techniques to provide estimated answers—for example, if out of 1000 simulated days there are 122 in which the pharmacist is still working at 5:30, we would estimate that the answer to question 2 is 0.122.

In order to be able to execute such an analysis, one must have some knowledge of probability so as to decide on certain probability distributions and questions such as whether appropriate random variables are to be assumed independent or not. A review of probability is provided in Chapter 2. The bases of a simulation study are so-called random numbers. A discussion of these quantities and how they are computer generated is presented in Chapter 3. Chapters 4 and 5 show how one can use random numbers to generate the values of random variables having arbitrary distributions. Discrete distributions are considered in Chapter 4 and continuous ones in Chapter 5. Chapter 6 introduces the multivariate normal distribution, and shows how to generate random variables having this joint distribution. Copulas, useful for modeling the joint distributions of random variables, are also introduced in Chapter 6. After completing Chapter 6, the reader should have some insight into the construction of a probability model for a given system and also how to use random numbers to generate the values of random quantities related to this model. The use of these generated values to track the system as it evolves continuously over time—that is, the actual simulation of the system—is discussed in Chapter 7, where we present the concept of "discrete events" and indicate how to utilize these entities to obtain a systematic approach to simulating systems. The discrete event simulation approach leads to a computer program, which can be written in whatever language the reader is comfortable in, that simulates the system a large number of times. Some hints concerning the verification of this program—to ascertain that it is actually doing what is desired—are also given in Chapter 7. The use of the outputs of a simulation study to answer probabilistic questions concerning the model necessitates the use of the theory of statistics, and this subject is introduced in Chapter 8. This chapter starts with the simplest and most basic concepts in statistics and continues toward "bootstrap statistics," which is quite useful in simulation. Our study of statistics indicates the importance of the variance of the estimators obtained from a simulation study as an indication of the efficiency of the simulation. In particular, the smaller this variance is, the smaller is the amount of simulation needed to obtain a fixed precision. As a result we are led, in Chapters 9 and 10, to ways of obtaining new estimators that are improvements over the raw simulation estimators because they have reduced variances. This topic of variance reduction is extremely important in a simulation study because it can substantially improve its efficiency. Chapter 11 shows how one can use the results of a simulation to verify, when some real-life data are available, the appropriateness of the probability model (which we have simulated) to the real-world situation. Chapter 12 introduces the important topic of Markov chain Monte Carlo methods. The use of these methods has, in recent years, greatly expanded the class of problems that can be attacked by simulation.

Exercises

1. The following data yield the arrival times and service times that each customer will require, for the first 13 customers at a single server system. Upon arrival, a customer either enters service if the server is free or joins the waiting line. When the server completes work on a customer, the next one in line (i.e., the one who has been waiting the longest) enters service.

Arrival Times:	12	31	63	95	99	154	198	221	304	346	411	455	537
Service Times:	40	32	55	48	18	50	47	18	28	54	40	72	12

(a) Determine the departure times of these 13 customers.
(b) Repeat (a) when there are two servers and a customer can be served by either one.
(c) Repeat (a) under the new assumption that when the server completes a service, the next customer to enter service is the one who has been waiting the least time.

2. Consider a service station where customers arrive and are served in their order of arrival. Let A_n, S_n, and D_n denote, respectively, the arrival time, the service time, and the departure time of customer n. Suppose there is a single server and that the system is initially empty of customers.

(a) With $D_0 = 0$, argue that for $n > 0$

$$D_n - S_n = \text{Maximum}\{A_n, D_{n-1}\}$$

(b) Determine the corresponding recursion formula when there are two servers.
(c) Determine the corresponding recursion formula when there are k servers.
(d) Write a computer program to determine the departure times as a function of the arrival and service times and use it to check your answers in parts (a) and (b) of Exercise 1.

Elements of Probability

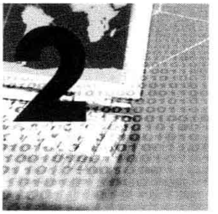

2.1 Sample Space and Events

Consider an experiment whose outcome is not known in advance. Let S, called the sample space of the experiment, denote the set of all possible outcomes. For example, if the experiment consists of the running of a race among the seven horses numbered 1 through 7, then

$$S = \{\text{all orderings of } (1, 2, 3, 4, 5, 6, 7)\}$$

The outcome $(3, 4, 1, 7, 6, 5, 2)$ means, for example, that the number 3 horse came in first, the number 4 horse came in second, and so on.

Any subset A of the sample space is known as an event. That is, an event is a set consisting of possible outcomes of the experiment. If the outcome of the experiment is contained in A, we say that A has occurred. For example, in the above, if

$$A = \{\text{all outcomes in } S \text{ starting with } 5\}$$

then A is the event that the number 5 horse comes in first.

For any two events A and B we define the new event $A \cup B$, called the union of A and B, to consist of all outcomes that are either in A or B or in both A and B. Similarly, we define the event AB, called the intersection of A and B, to consist of all outcomes that are in both A and B. That is, the event $A \cup B$ occurs if either A or B occurs, whereas the event AB occurs if both A and B occur. We can also define unions and intersections of more than two events. In particular, the union of the events $A_1, \ldots, A_n$—designated by $\cup_{i=1}^{n} A_i$—is defined to consist of all outcomes that are in any of the A_i. Similarly, the intersection of the events $A_1, \ldots, A_n$—designated by $A_1 A_2 \cdots A_n$—is defined to consist of all outcomes that are in all of the A_i.

Simulation. DOI: http://dx.doi.org/10.1016/B978-0-12-415825-2.00002-4

For any event A we define the event A^c, referred to as the complement of A, to consist of all outcomes in the sample space S that are not in A. That is, A^c occurs if and only if A does not. Since the outcome of the experiment must lie in the sample space S, it follows that S^c does not contain any outcomes and thus cannot occur. We call S^c the null set and designate it by $\emptyset$. If $AB = \emptyset$ so that A and B cannot both occur (since there are no outcomes that are in both A and B), we say that A and B are mutually exclusive.

2.2 Axioms of Probability

Suppose that for each event A of an experiment having sample space S there is a number, denoted by $P(A)$ and called the probability of the event A, which is in accord with the following three axioms:

Axiom 1 $0 \leqslant P(A) \leqslant 1$

Axiom 2 $P(S) = 1$

Axiom 3 *For any sequence of mutually exclusive events* $A_1, A_2, \ldots$

$$P\left(\bigcup_{i=1}^{n} A_i\right) = \sum_{i=1}^{n} P(A_i), \quad n = 1, 2, \ldots, \infty.$$

Thus, Axiom 1 states that the probability that the outcome of the experiment lies within A is some number between 0 and 1; Axiom 2 states that with probability 1 this outcome is a member of the sample space; and Axiom 3 states that for any set of mutually exclusive events, the probability that at least one of these events occurs is equal to the sum of their respective probabilities.

These three axioms can be used to prove a variety of results about probabilities. For instance, since A and A^c are always mutually exclusive, and since $A \cup A^c = S$, we have from Axioms 2 and 3 that

$$1 = P(S) = P(A \cup A^c) = P(A) + P(A^c)$$

or equivalently

$$P(A^c) = 1 - P(A)$$

In words, the probability that an event does not occur is 1 minus the probability that it does.