

CHAPTER 1

INTRODUCTION

1.1 WHAT IS REGRESSION ANALYSIS?

Regression analysis is a conceptually simple method for investigating functional relationships among variables. A real estate appraiser may wish to relate the sale price of a home from selected physical characteristics of the building and taxes (local, school, county) paid on the building. We may wish to examine whether cigarette consumption is related to various socioeconomic and demographic variables such as age, education, income, and price of cigarettes. The relationship is expressed in the form of an equation or a model connecting the *response* or *dependent* variable and one or more *explanatory* or *predictor* variables. In the cigarette consumption example, the response variable is cigarette consumption (measured by the number of packs of cigarette sold in a given state on a per capita basis during a given year) and the explanatory or predictor variables are the various socioeconomic and demographic variables. In the real estate appraisal example, the response variable is the price of a home and the explanatory or predictor variables are the characteristics of the building and taxes paid on the building.

We denote the response variable by Y and the set of predictor variables by X_1, X_2, \dots, X_p , where p denotes the number of predictor variables. The true relationship between Y and X_1, X_2, \dots, X_p can be approximated by the regression model

$$Y = f(X_1, X_2, \dots, X_p) + \varepsilon, \quad (1.1)$$

where ε is assumed to be a random error representing the discrepancy in the approximation. It accounts for the failure of the model to fit the data exactly. The function $f(X_1, X_2, \dots, X_p)$ describes the relationship between Y and X_1, X_2, \dots, X_p . An example is the linear regression model

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_p X_p + \varepsilon, \quad (1.2)$$

where $\beta_0, \beta_1, \dots, \beta_p$, called the *regression parameters* or *coefficients*, are unknown constants to be determined (estimated) from the data. We follow the commonly used notational convention of denoting unknown parameters by Greek letters.

The predictor or explanatory variables are also called by other names such as *independent variables*, *covariates*, *regressors*, *factors*, and *carriers*. The name independent variable, though commonly used, is the least preferred, because in practice the predictor variables are rarely independent of each other.

1.2 PUBLICLY AVAILABLE DATA SETS

Regression analysis has numerous areas of applications. A partial list would include economics, finance, business, law, meteorology, medicine, biology, chemistry, engineering, physics, education, sports, history, sociology, and psychology. A few examples of such applications are given in Section 1.3. Regression analysis is learned most effectively by analyzing data that are of direct interest to the reader. We invite the readers to think about questions (in their own areas of work, research, or interest) that can be addressed using regression analysis. Readers should collect the relevant data and then apply the regression analysis techniques presented in this book to their own data. To help the reader locate real-life data, this section provides some sources and links to a wealth of data sets that are available for public use.

A number of data sets are available in books and on the Internet. The book by Hand et al. (1994) contains data sets from many fields. These data sets are small in size and are suitable for use as exercises. The book by Chatterjee et al. (1995) provides numerous data sets from diverse fields. The data are

included in a diskette that comes with the book and can also be found at the Website.¹

Data sets are also available on the Internet at many other sites. Some of the Websites given below allow the direct copying and pasting into the statistical package of choice, while others require downloading the data file and then importing them into a statistical package. Some of these sites also contain further links to yet other data sets or statistics-related Websites.

The Data and Story Library (DASL, pronounced “dazzle”) is one of the most interesting sites that contains a number of data sets accompanied by the “story” or background associated with each data set. DASL is an online library² of data files and stories that illustrate the use of basic statistical methods. The data sets cover a wide variety of topics. DASL comes with a powerful search engine to locate the story or data file of interest.

Another Website, which also contains data sets arranged by the method used in the analysis, is the Electronic Dataset Service.³ The site also contains many links to other data sources on the Internet.

Finally, this book has a Website,⁴ which contains, among other things, all the data sets that are included in this book and more. These and other data sets can be found at the Book’s Website.

1.3 SELECTED APPLICATIONS OF REGRESSION ANALYSIS

Regression analysis is one of the most widely used statistical tools because it provides simple methods for establishing a functional relationship among variables. It has extensive applications in many subject areas. The cigarette consumption and the real estate appraisal, mentioned above, are but two examples. In this section, we give a few additional examples demonstrating the wide applicability of regression analysis in real-life situations. Some of the data sets described here will be used later in the book to illustrate regression techniques or in the exercises at the end of various chapters.

1.3.1 Agricultural Sciences

The Dairy Herd Improvement Cooperative (DHI) in upstate New York collects and analyzes data on milk production. One question of interest here

¹ <http://www.stern.nyu.edu/~jsimonof/Casebook>

² <http://lib.stat.cmu.edu/DASL>

³ <http://www-unix.oit.umass.edu/~statdata>

⁴ <http://www.aucegypt.edu/faculty/hadi/RABE6>

Table 1.1 Variables in Milk Production Data

Variable	Definition
Current	Current month milk production in pounds
Previous	Previous month milk production in pounds
Fat	Percent of fat in milk
Protein	Percent of protein in milk
Days	Number of days since present lactation
Lactation	Number of lactations
I79	Indicator variable (0 if Days \leq 79 and 1 if Days $>$ 79)

is how to develop a suitable model to predict current milk production from a set of measured variables. The response variable (current milk production in pounds) and the predictor variables are given in Table 1.1. Samples are taken once a month during milking. The period that a cow gives milk is called lactation. Number of lactations is the number of times a cow has calved or given milk. The recommended management practice is to have the cow produce milk for about 305 days and then allow a 60-day rest period before beginning the next lactation. The data set, consisting of 199 observations, was compiled from the DHI milk production records. The Milk Production data can be found at the Book's Website.

1.3.2 Industrial and Labor Relations

In 1947, the United States Congress passed the Taft–Hartley Amendments to the Wagner Act. The original Wagner Act had permitted the unions to use a *Closed Shop Contract*⁵ unless prohibited by state law. The Taft–Hartley Amendments made the use of Closed Shop Contract illegal and gave individual states the right to prohibit union shops⁶ as well. These right-to-work laws have caused a wave of concern throughout the labor movement. A question of interest here is: What are the effects of these laws on the cost of living for a four-person family living on an intermediate budget in the United States? To answer this question a data set consisting of 38 geographic locations

⁵ Under a Closed Shop Contract provision, all employees must be union members at the time of hire and must remain members as a condition of employment.

⁶ Under a Union Shop clause, employees are not required to be union members at the time of hire, but must become a member within two months, thus allowing the employer complete discretion in hiring decisions.

Table 1.2 Variables in Right-To-Work Laws Data

Variable	Definition
COL	Cost of living for a four-person family
PD	Population density (person per square mile)
URate	State unionization rate in 1978
Pop	Population in 1975
Taxes	Property taxes in 1972
Income	Per capita income in 1974
RTWL	Indicator variable (1 if there are right-to-work laws in the state and 0 otherwise)

has been assembled from various sources. The variables used are defined in Table 1.2. The Right-To-Work Laws data can be found at the Book's Website.

1.3.3 Government

Information about domestic immigration (the movement of people from one state or area of a country to another) is important to state and local governments. It is of interest to build a model that predicts domestic immigration or to answer the question of why do people leave one place to go to another? There are many factors that influence domestic immigration, such as weather conditions, crime, taxes, and unemployment rates. A data set for the 48 contiguous states has been created. Alaska and Hawaii are excluded from the analysis because the environments of these states are significantly different from the other 48, and their locations present certain barriers to immigration. The response variable here is net domestic immigration, which represents the net movement of people into and out of a state over the period 1990–1994 divided by the population of the state. Eleven predictor variables thought to influence domestic immigration are defined in Table 1.3 and can be found at the Book's Website.

1.3.4 History

A question of historical interest is how to estimate the age of historical objects based on some age-related characteristics of the objects. For example, the variables in Table 1.4 can be used to estimate the age of Egyptian skulls.

Table 1.3 Variables in Study of Domestic Immigration

Variable	Definition
State	State name
NDIR	Net domestic immigration rate over the period 1990–1994
Unemp	Unemployment rate in the civilian labor force in 1994
Wage	Average hourly earnings of production workers in manufacturing in 1994
Crime	Violent crime rate per 100,000 people in 1993
Income	Median household income in 1994
Metrop	Percentage of state population living in metropolitan areas in 1992
Poor	Percentage of population who fall below the poverty level in 1994
Taxes	Total state and local taxes per capita in 1993
Educ	Percentage of population 25 years or older who have a high school degree or higher in 1990
BusFail	The number of business failures divided by the population of the state in 1993
Temp	Average of the 12 monthly average temperatures (in degrees Fahrenheit) for the state in 1993
Region	Region in which the state is located (northeast, south, midwest, west)

Table 1.4 Variables in Egyptian Skulls Data

Variable	Definition
Year	Approximate year of skull formation (negative = B.C.; positive = A.D.)
MB	Maximum breadth of skull
BH	Basibregmatic height of skull
BL	Basialveolar length of skull
NH	Nasal Height of skull

Here the response variable is Year and the other four variables are possible predictors. There are 150 observations in this data set. The original source of the data is Thomson and Randall-Maciver (1905), but they can be found in Hand et al. (1994, pp. 299–301). An analysis of the data can be found in Manly (1986). The Egyptian Skulls data can be found at the Book’s Website.

Table 1.5 Variables in Study of Water Pollution in New York Rivers

Variable	Definition
Y	Mean nitrogen concentration (mg/liter) based on samples taken at regular intervals during the spring, summer, and fall months
X_1	Agriculture: percentage of land area currently in agricultural use
X_2	Forest: percentage of forest land
X_3	Residential: percentage of land area in residential use
X_4	Commercial/Industrial: percentage of land area in either commercial or industrial use

1.3.5 Environmental Sciences

In a 1976 study exploring the relationship between water quality and land use, Haith (1976) obtained the measurements (shown in Table 1.5) on 20 river basins in New York State. A question of interest here is how the land use around a river basin contributes to the water pollution as measured by the mean nitrogen concentration (mg/liter). The dataset can be found at the Book's Website.

1.3.6 Industrial Production

Nambe Mills in Santa Fe, New Mexico, makes a line of tableware made from sand casting a special alloy of metals. After casting, the pieces go through a series of shaping, grinding, buffing, and polishing steps. Data was collected for 59 items produced by the company. The relation between the polishing time and the product diameters and the product types (Bowl, Casserole, Dish, Tray, and Plate) are used to estimate the polishing time for new products which are designed or suggested for design and manufacture. The variables representing product types are coded as binary variables (1 corresponds to the type and 0 otherwise). Diam is the diameter of the item (in inches), polishing time is measured in minutes, and price in dollars. The polishing time is the major item in the cost of the product. The production decision will be based on the estimated time of polishing. The data is obtained from the DASL library, can be found there and also in the file `Industrial.Production.csv` at the Book's Website.

1.3.7 The Space Shuttle Challenger

The explosion of the space shuttle Challenger in 1986 killing the crew was a shattering tragedy. To look into the case a Presidential Commission was appointed. The O-rings in the booster rockets, which are used in space launching, play a very important part in its safety. The rigidity of the O-rings is thought to be affected by the temperature at launching. There are six O-rings in a booster rocket. The data consists of two variables: the number of rings damaged and the temperature at launchings of the 23 flights. The data set can be found at the Book's Website. The analysis performed before the launch did not include the launches in which no O-ring was damaged and came to the wrong conclusion. A detailed discussion of the problem is found in *The Flight of the Space Shuttle Challenger* in Chatterjee et al. (1995, pp. 33–35). Note here that the response variable is a proportion bounded between 0 and 1.

1.3.8 Cost of Health Care

The cost of delivery of health care has become an important concern. Getting data on this topic is extremely difficult because it is highly proprietary. These data were collected by the Department of Health and Social Services of the State of New Mexico and cover 52 of the 60 licensed facilities in New Mexico in 1988. The variables in these data are the characteristics which describe the facilities size, volume of usage, expenditures, and revenue. The location of the facility is also indicated, whether it is in the rural or nonrural area. Specific definitions of the variables are given in Table 1.6 and the data can be found at the Book's Website. There are several ways of looking at a

Table 1.6 Variables in Cost of Health Care Data

Variable	Definition
RURAL	Rural home (1) and nonrural home (0)
BED	Number of beds in home
MCDAYS	Annual medical in-patient days (hundreds)
TDAYS	Annual total patient days (hundreds)
PCREV	Annual total patient care revenue (\$100)
NSAL	Annual nursing salaries (\$100)
FEXP	Annual facilities expenditures (\$100)
NETREV	$PCREV - NSAL - FEXP$

body of data and extracting various kinds of information. For example, (a) Are rural facilities different from nonrural facilities? and (b) How do the hospital characteristics affect the total patient care revenue?

1.4 STEPS IN REGRESSION ANALYSIS

Regression analysis includes the following steps:

- Statement of the problem
- Selection of potentially relevant variables
- Data collection
- Model specification
- Choice of fitting method
- Model fitting
- Model validation and criticism
- Using the chosen model(s) for the solution of the posed problem.

These steps are examined below.

1.4.1 Statement of the Problem

Regression analysis usually starts with a formulation of the problem. This includes the determination of the question(s) to be addressed by the analysis. The problem statement is the first and perhaps the most important step in regression analysis. It is important because an ill-defined problem or a misformulated question can lead to wasted effort. It can lead to the selection of irrelevant set of variables or to a wrong choice of the statistical method of analysis. A question that is not carefully formulated can also lead to the wrong choice of a model. Suppose we wish to determine whether or not an employer is discriminating against a given group of employees, say women. Data on salary, qualifications, and gender are available from the company's record to address the issue of discrimination. There are several definitions of employment discrimination in the literature. For example, discrimination occurs when on the average (a) women are paid less than equally qualified men, or (b) women are more qualified than equally paid men. To answer the question: "On the average, are women paid less than equally qualified men?"

we choose salary as a response variable, and qualification and gender as predictor variables. But to answer the question: “On the average, are women more qualified than equally paid men?” we choose qualification as a response variable and salary and gender as predictor variables, that is, the roles of variables have been switched.

1.4.2 Selection of Potentially Relevant Variables

The next step after the statement of the problem is to select a set of variables that are thought by the experts in the area of study to explain or predict the response variable. The response variable is denoted by Y and the explanatory or predictor variables are denoted by X_1, X_2, \dots, X_p , where p denotes the number of predictor variables. An example of a response variable is the price of a single-family house in a given geographical area. A possible relevant set of predictor variables in this case is: area of the lot, area of the house, age of the house, number of bedrooms, number of bathrooms, type of neighborhood, style of the house, amount of real estate taxes, and so forth.

1.4.3 Data Collection

The next step after the selection of potentially relevant variables is to collect the data from the environment under study to be used in the analysis. Sometimes the data are collected in a controlled setting so that factors that are not of primary interest can be held constant. More often the data are collected under nonexperimental conditions where very little can be controlled by the investigator. In either case, the collected data consist of observations on n subjects. Each of these n observations consists of measurements for each of the potentially relevant variables. The data are usually recorded as in Table 1.7. A column in Table 1.7 represents a variable, whereas a row represents an observation, which is a set of $p + 1$ values for a single subject (e.g., a house); one value for the response variable and one value for each of the p predictors. The notation x_{ij} refers to the i th value of the j th variable. The first subscript refers to observation number and the second refers to variable number.

Each of the variables in Table 1.7 can be classified as either *quantitative* or *qualitative*. Examples of quantitative variables are the house price, number of bedrooms, age, and taxes. Examples of qualitative variables are neighborhood type (e.g., good or bad neighborhood) and house style (e.g., ranch, colonial, etc.). In this book we deal mainly with the cases where the response variable is

Table 1.7 Notation for Data Used in Regression Analysis

Observation Number	Response Variable	Predictors			
	Y	X_1	X_2	\dots	X_p
1	y_1	x_{11}	x_{12}	\dots	x_{1p}
2	y_2	x_{21}	x_{22}	\dots	x_{2p}
3	y_3	x_{31}	x_{32}	\dots	x_{3p}
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
n	y_n	x_{n1}	x_{n2}	\dots	x_{np}

quantitative. A technique used in cases where the response variable is *binary*⁷ is called *logistic regression*. This is introduced in Chapter 13. In regression analysis, the predictor variables can be either quantitative and/or qualitative. For the purpose of computations, however, the qualitative variables, if any, have to be coded into a set of *indicator* or *dummy* variables as discussed in Chapter 6.

If all predictor variables are qualitative, the techniques used in the analysis of the data are called the *analysis of variance* techniques. Although the analysis of variance techniques can be introduced and explained as methods in their own right,⁸ it is shown in Chapter 6 that they are special cases of regression analysis. If some of the predictor variables are quantitative while others are qualitative, regression analysis in these cases is called the *analysis of covariance*.

1.4.4 Model Specification

The form of the model that is thought to relate the response variable to the set of predictor variables can be specified initially by the experts in the area of study based on their knowledge or their objective and/or subjective judgments. The hypothesized model can then be either confirmed or refuted by the analysis of the collected data. Note that the model needs to be specified only in form, but it can still depend on unknown parameters. We need to select the form of the function $f(X_1, X_2, \dots, X_p)$ in (1.1). This function can

⁷ A variable that can take only one of two possible values such as yes or no, 1 or 0, and success or failure, is called a binary variable.

⁸ See, for example, the books by Scheffé (1959), Iversen (1976), Wildt and Ahtola (1978), Krishnaiah (1980), Iversen and Norpoth (1987), Lindman (1992), and Christensen (1996).

be classified into two types: *linear* and *nonlinear*. An example of a linear function is

$$Y = \beta_0 + \beta_1 X_1 + \varepsilon \quad (1.3)$$

while a nonlinear function is

$$Y = \beta_0 + e^{\beta_1 X_1} + \varepsilon. \quad (1.4)$$

Note that the term *linear* (*nonlinear*) here does not describe the relationship between Y and X_1, X_2, \dots, X_p . It is related to the fact that the regression parameters enter the equation linearly (nonlinearly). Each of the following models is linear:

$$Y = \beta_0 + \beta_1 X + \beta_2 X^2 + \varepsilon,$$

$$Y = \beta_0 + \beta_1 \ln X + \varepsilon,$$

because in each case the parameters enter linearly although the relationship between Y and X is nonlinear. This can be seen if the two models are reexpressed, respectively, as follows:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \varepsilon,$$

$$Y = \beta_0 + \beta_1 X_1 + \varepsilon,$$

where in the first equation we have $X_1 = X$ and $X_2 = X^2$ and in the second equation we have $X_1 = \ln X$. The variables here are *reexpressed* or *transformed*. Transformation is dealt with in Chapter 7. All nonlinear functions that can be transformed into linear functions are called *linearizable* functions. Accordingly, the class of linear models is actually wider than it might appear at first sight because it includes all linearizable functions. Note, however, that not all nonlinear functions are linearizable. For example, it is not possible to linearize the nonlinear function in (1.4). Some authors refer to nonlinear functions that are not linearizable as *intrinsically* nonlinear functions.

A regression equation containing only one predictor variable is called a *simple regression equation*. An equation containing more than one predictor variable is called a *multiple regression equation*. An example of simple regression would be an analysis in which the time to repair a machine is studied in relation to the number of components to be repaired. Here we have one response variable (time to repair the machine) and one predictor variable (number of components to be repaired). An example of a very complex multiple regression situation would be an attempt to explain the age-adjusted

mortality rates prevailing in different geographic regions (response variable) by a large number of environmental and socioeconomic factors (predictor variables). Both types of problems are treated in this book. These two particular examples are studied, one in Chapter 3, the other in Chapter 12.

In certain applications the response variable can actually be a set of variables, Y_1, Y_2, \dots, Y_q , say, which are thought to be related to the same set of predictor variables, X_1, X_2, \dots, X_p . For example, Bartlett et al. (1998) present a data set on 148 healthy people. Eleven variables are measured; six variables represent different types of measured sensory thresholds (e.g., vibration, hand and foot temperatures) and five a priori selected baseline covariates (e.g., age, gender, height, and weight) that may have systematic effects on some or all of the six sensory thresholds. Here we have six response variables and five predictor variables. This data set, which we refer to as the QST (*quantitative sensory testing*) data, is not listed here due to its size (148 observations) but it can be found at the Book's Website. For further description of the data and objectives of the study, see Bartlett et al. (1998).

When we deal only with one response variable, regression analysis is called *univariate* regression and in cases where we have two or more response variables, the regression is called *multivariate* regression. Simple and multiple regressions should not be confused with univariate versus multivariate regressions. The distinction between simple and multiple regressions is determined by the number of predictor variables (simple means one predictor variable and multiple means two or more predictor variables), whereas the distinction between univariate and multivariate regressions is determined by the number of response variables (univariate means one response variable and multivariate means two or more response variables). In this book we consider only univariate regression (both simple and multiple, linear and nonlinear). Multivariate regression is treated in books on multivariate analysis such as Rencher (1995), Johnson and Wichern (1992), and Johnson (1998). In this book the term regression will be used to mean univariate regression.

The various classifications of regression analysis we discussed above are shown in Table 1.8.

1.4.5 Method of Fitting

After the model has been defined and the data have been collected, the next task is to estimate the parameters of the model based on the collected data. This is also referred to as *parameter estimation* or *model fitting*. The most commonly used method of estimation is called the *least squares* method. Under certain assumptions (to be discussed in detail in this book), least

Table 1.8 Various Classifications of Regression Analysis

Type of Regression	Conditions
Univariate	Only one quantitative response variable
Multivariate	Two or more quantitative response variables
Simple	Only one predictor variable
Multiple	Two or more predictor variables
Linear	All parameters enter the equation linearly, possibly after transformation of the data
Nonlinear	The relationship between the response and some of the predictors is nonlinear or some of the parameters appear nonlinearly, but no transformation is possible to make the parameters appear linearly
Analysis of variance	All predictors are qualitative variables
Analysis of covariance	Some predictors are quantitative variables and others are qualitative variables
Logistic	The response variable is qualitative

squares method produce estimators with desirable properties. In this book we will deal mainly with the least squares method and its variants (e.g., weighted least squares). In some instances (e.g., when one or more of the assumptions does not hold) other estimation methods may be superior to least squares. The other estimation methods that we consider in this book are the *maximum likelihood* method, the *ridge regression*, and the *principal components* method.

1.4.6 Model Fitting

The next step in the analysis is to estimate the regression parameters or to fit the model to the collected data using the chosen estimation method (e.g., least squares). The estimates of the regression parameters $\beta_0, \beta_1, \dots, \beta_p$ in (1.1) are denoted by $\hat{\beta}_0, \hat{\beta}_1, \dots, \hat{\beta}_p$. The estimated regression equation then becomes

$$\hat{Y} = \hat{\beta}_0 + \hat{\beta}_1 X_1 + \hat{\beta}_2 X_2 + \cdots + \hat{\beta}_p X_p. \quad (1.5)$$

A *hat* on top of a parameter denotes an estimate of the parameter. The value \hat{Y} (pronounced as *Y-hat*) is called the *fitted* value. Using (1.5), we can compute n fitted values, one for each of the n observations in our data. For example,

the i th fitted value \hat{y}_i is

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 x_{i1} + \hat{\beta}_2 x_{i2} + \cdots + \hat{\beta}_p x_{ip}, \quad i = 1, 2, \dots, n, \quad (1.6)$$

where x_{i1}, \dots, x_{ip} are the values of the p predictor variables for the i th observation.

Note that (1.5) can be used to predict the response variable for any values of the predictor variables not observed in our data. In this case, the obtained \hat{Y} is called the *predicted* value. The difference between fitted and predicted values is that the fitted value refers to the case where the values used for the predictor variables correspond to one of the n observations in our data, but the predicted values are obtained for any set of values of the predictor variables. It is generally not recommended to predict the response variable for a set of values of the predictor variables far outside the range of our data. In cases where the values of the predictor variables represent future values of the predictors, the predicted value is referred to as the *forecasted* value.

1.4.7 Model Criticism and Selection

The validity of a statistical method, such as regression analysis, depends on certain assumptions. Assumptions are usually made about the data and the model. The accuracy of the analysis and the conclusions derived from an analysis depend crucially on the validity of these assumptions. Before using (1.5) for any purpose, we first need to determine whether the specified assumptions hold. We need to address the following questions:

1. What are the required assumptions?
2. For each of these assumptions, how do we determine whether or not the assumption is valid?
3. What can be done in cases where one or more of the assumptions does not hold?

The standard regression assumptions will be specified and the above questions will be addressed in great detail in various parts of this book. We emphasize here that validation of the assumptions must be made *before* any conclusions are drawn from the analysis. Regression analysis is viewed here as an *iterative* process, a process in which the outputs are used to diagnose, validate, criticize, and possibly modify the inputs. The process has to be repeated until a satisfactory output has been obtained. A satisfactory output is an estimated model that satisfies the assumptions and fits the data reasonably

well. This iterative process is illustrated schematically in Figure 1.1. This dynamic iterative regression process can be implemented with the aid of the flowchart in Figure 1.2. This process will be illustrated by examples in later chapters.

1.4.8 Objectives of Regression Analysis

The explicit determination of the regression equation is the most important product of the analysis. It is a summary of the relationship between Y (the response variable) and the set of predictor variables X_1, X_2, \dots, X_p . The equation may be used for several purposes. It may be used to evaluate the importance of individual predictors, to analyze the effects of policy that involves changing values of the predictor variables, or to forecast values of the response variable for a given set of predictors. Although the regression equation is the final product, there are many important by-products. We view regression analysis as a set of data analytic techniques that are used to help understand the interrelationships among variables in a certain environment. The task of regression analysis is to learn as much as possible about the environment reflected by the data. We emphasize that what is uncovered along the way to the formulation of the equation may often be as valuable and informative as the final equation.

1.5 SCOPE AND ORGANIZATION OF THE BOOK

This book can be used by all who analyze data. A knowledge of matrix algebra is not necessary. We have seen excellent regression analysis done by people who have no knowledge of matrix theory. A knowledge of matrix algebra is certainly very helpful in understanding the theory. We have provided appendices which use matrix algebra for readers who are familiar with that topic. Matrix algebra permits expression of regression results much more compactly and is essential for the mathematical derivation of the results.

Lack of knowledge of matrix algebra should not deter anyone from using this book and doing regression analysis. For readers who are not familiar with matrix algebra but who wish to benefit from the material in the appendices, we recommend reading the relatively short book by Hadi (1996), *Matrix Algebra as a Tool*. We believe that the majority, if not all, of our readers can read it entirely on their own or with minimal assistance.

There are no formal derivations in the text and readers interested in mathematical derivations are referred to a number of books that contain




Figure 1.1 A schematic illustration of the iterative nature of the regression process.

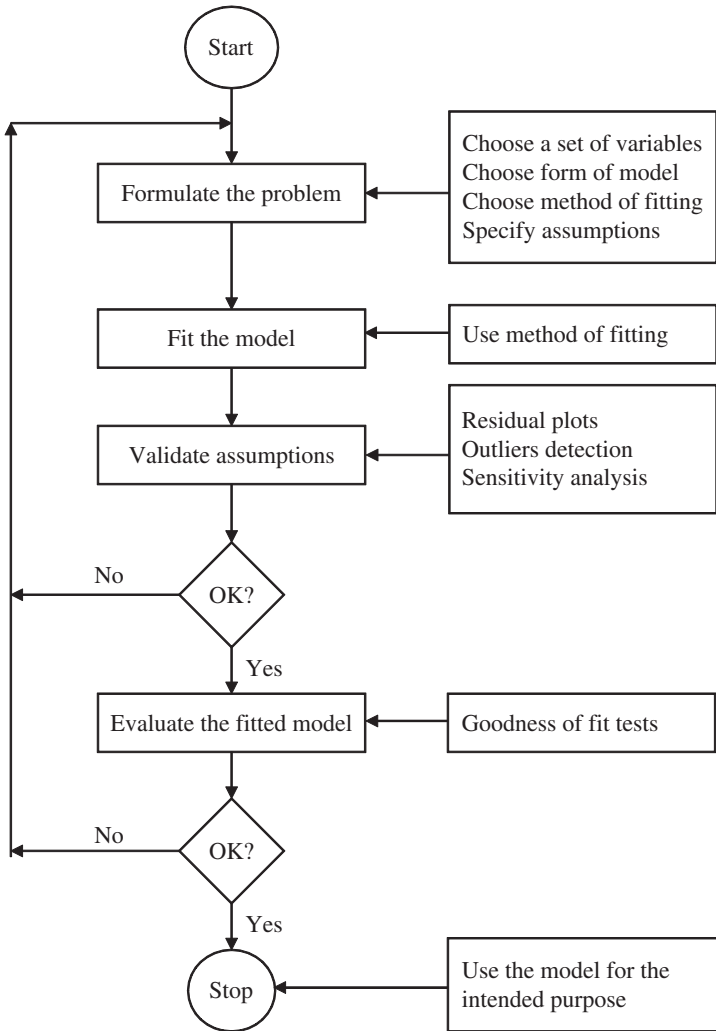


Figure 1.2 A flowchart illustrating the dynamic iterative regression process.

formal derivations of the regression formulas. Formulas are presented, but only for purposes of reference.

The rest of the book is organized as follows: Chapter 2 gives a brief but sufficient introduction to the R language with no prior knowledge of R. Then, we begin with the simple linear regression model in Chapter 3. The simple regression model is then extended to the multiple regression model in Chapter 4. In both chapters, the model is formulated, assumptions are specified, and the key theoretical results are stated and illustrated by examples. For simplicity of presentation and for pedagogical reasons, the analysis and conclusions in Chapters 3 and 4 are made under the presumption that the standard regression assumptions are valid. Chapter 5 addresses the issue of assumptions validation and the detection and correction of model violations.

Each of the Chapters 6–12 deals with a special regression problem. Chapter 6 deals with the case where some or all of the predictor variables are qualitative. Chapter 7 deals with data transformation. Chapter 8 presents situations where a variant of the least squares method is needed. This method is called the *weighted least squares* method. Chapter 9 discusses the problem that arises when the observations are correlated. This problem is known as the *autocorrelation* problem. Chapters 10 and 11 presents methods for the detection and correction of an important problem called *collinearity*. Collinearity occurs when the predictor variables are highly correlated.

Chapter 12 presents variable selection methods – computer methods for selecting the best and most parsimonious model(s). Before applying any of the variable selection methods, we assume in this chapter that questions of assumptions validation and model violations have already been addressed and settled satisfactorily.

Chapters 3–12 dealt with the case where the response variable is quantitative. Chapter 13 discusses logistic regression, the method used when the response variable is categorical. Logistic regression is studied because it is an important tool with many applications. Beside binary logistic regression, we have now included a discussion of multinomial logistic regression. This extends the application of logistic regression to more diverse situations. The categories in some multinomial are ordered, for example, in attitude surveys. We also discuss the application of the logistic model to ordered response variables.

The book concludes with Chapter 14 entitled Further Topics. Two topics are discussed in this chapter. One extends the concept of linear models so that regression and logistic models are all viewed as special cases of the linear model. This extends the range of applications of linear models to more

diverse situations. We also discuss Poisson regression, often used to model count data. A brief discussion of robust regression with illustrative examples is also given in this chapter.

We recommend that the chapters be covered in the same sequence as they are presented, although Chapters 6–13 can be covered in any order after Chapter 5, as long as Chapter 10 is covered before Chapter 11, and Chapter 8 is covered before Chapters 13 and 14.

EXERCISES

1.1 Classify each of the following variables as either quantitative or qualitative. If a variable is qualitative, state the possible categories.

- | | |
|-------------------------|------------------------------------|
| (a) Geographical region | (b) Number of children in a family |
| (c) Price of a house | (d) Race |
| (e) Temperature | (f) Fuel consumption |
| (g) Employment rate | (h) Political party preference |

1.2 Give two examples in any area of interest to you (other than those already presented in this chapter) where regression analysis can be used as a data analytic tool to answer some questions of interest. For each example:

- What is the question of interest?
- Identify the response and the predictor variables.
- Classify each of the variables as either quantitative or qualitative.
- Which type of regression (see Table 1.8) can be used to analyze the data?
- Give a possible form of the model and identify its parameters.

1.3 In each of the following sets of variables, identify which of the variables can be regarded as a response variable and which can be used as predictors? (Explain)

- Number of cylinders and gasoline consumption of cars
- SAT scores, grade point average, and college admission
- Supply and demand of certain goods
- Company's assets, return on a stock, and net sales
- The distance of a race, the time to run the race, and the weather conditions at the time of running

- (f) The weight of a person, whether or not the person is a smoker, and whether or not the person has a lung cancer
- (g) The height and weight of a child, his/her parents' height and weight, and the gender and age of the child.

1.4 For each of the sets of variables in Exercise 1.3:

- (a) Classify each variable as either quantitative or qualitative.
- (b) Which type of regression (see Table 1.8) can be used in the analysis of the data?