1

Introduction and Mathematical Foundations

LEARNING OUTCOMES

In this chapter, you will learn how to

- Describe the key steps involved in building an econometric model
- Work with powers, exponents and logarithms
- Plot, interpret and calculate the roots of functions
- Use sigma (Σ) and pi (Π) notation
- Apply rules to differentiate various types functions
- Work with matrices
- Calculate the trace, inverse and eigenvalues of a matrix
- Construct and interpret utility functions

Learning econometrics is in many ways like learning a new language. To begin with, nothing makes sense and it is as if it is impossible to see through the fog created by all the unfamiliar terminology. While the way of writing the models – the *notation* – may make the situation appear more complex, in fact it is supposed to achieve the exact opposite. The ideas themselves are mostly not so complicated, it is just a matter of learning enough of the language that everything fits into place. So if you have never studied the subject before, then persevere through this preliminary chapter and you will hopefully be on your way to being fully fluent in econometrics!

This chapter comprises two parts. The first sets the scene for the book by discussing in broad terms the questions of what econometrics is, and the kinds of problems that can be tackled using econometrics. The second part of the chapter covers the mathematical techniques that underpin approaches to modelling and dealing with data in finance. Those with some prior background in algebra and introductory mathematics may skip the second part of this chapter without loss of continuity, but hopefully the material will also constitute a useful refresher for those who have studied mathematics but a long time ago!

1.1 What is Econometrics?

The literal meaning of the word 'econometrics' is 'measurement in economics'. The first five letters of the word suggest correctly that the origins of econometrics are rooted in economics. However, the main techniques employed for studying economic problems are of equal importance in financial applications. As the term is used in this book, financial econometrics will be defined as the *application of statistical techniques to problems in finance*. Financial econometrics can be useful for testing theories in finance, determining asset prices or returns, testing hypotheses concerning the relationships between variables, examining the effect on financial markets of changes in economic conditions, forecasting future values of financial variables and for financial decision-making. A list of possible examples of where econometrics may be useful is given in Box 1.1.

The list in Box 1.1 is of course by no means exhaustive, but it hopefully gives some flavour of the usefulness of econometric tools in terms of their financial applicability.

BOX 1.1 Examples of the uses of econometrics

(1) Testing whether financial markets are weak-form informationally efficient
(2) Testing whether the capital asset pricing model (CAPM) or arbitrage pricing theory (APT) represent superior models for the determination of returns on risky assets
(3) Measuring and forecasting the volatility of bond returns
(4) Explaining the determinants of bond credit ratings used by the ratings agencies
(5) Modelling long-term relationships between prices and exchange rates
(6) Determining the optimal hedge ratio for a spot position in oil
(7) Testing technical trading rules to determine which makes the most money
(8) Testing the hypothesis that earnings or dividend announcements have no effect on stock prices
(9) Testing whether spot or futures markets react more rapidly to news
(10) Forecasting the correlation between the stock indices of two countries.

1.2 Is Financial Econometrics Different from 'Economic Econometrics'?

As previously stated, the tools commonly used in financial applications are fundamentally the same as those used in economic applications, although the emphasis and the sets of problems that are likely to be encountered when analysing the two sets of data are somewhat different. Financial data often differ from macroeconomic data in terms of their frequency, accuracy, seasonality and other properties.

In economics, a serious problem is often a *lack of data at hand* for testing the theory or hypothesis of interest – this is sometimes called a 'small samples problem'. It might be, for example, that data are required on government budget deficits, or population figures, which are measured only on an annual basis. If the methods used to measure these quantities changed a quarter of a century ago, then only at most twenty-five of these annual observations are usefully available.

Two other problems that are often encountered in conducting applied econometric work in the arena of economics are those of *measurement error* and *data revisions*. These difficulties are simply that the data may be estimated, or measured with error, and will often be subject to several vintages of subsequent revisions. For example, a researcher may estimate an economic model of the effect on national output of investment in computer technology using a set of published data, only to find that the data for the last two years have been revised substantially in the next, updated publication.

These issues are usually of less concern in finance. Financial data come in many shapes and forms, but in general the prices and other entities that are recorded are those at which trades *actually took place*, or which were *quoted* on the screens of information providers. There exists, of course, the possibility for typos or for the data measurement method to change (for example, owing to stock index re-balancing or re-basing). But in general the measurement error and revisions problems are far less serious in the financial context.

Similarly, some sets of financial data are observed at much *higher frequencies* than macroeconomic data. Asset prices or yields are often available at daily, hourly or minute-by-minute frequencies. Thus the number of observations available for analysis can potentially be very large – perhaps thousands or even millions, making financial data the envy of macro-econometricians! The implication is that more powerful techniques can often be applied to financial than economic data, and that researchers may also have more confidence in the results.

Furthermore, the analysis of financial data also brings with it a number of new problems. While the difficulties associated with handling and processing such a large amount of data are not usually an issue given recent and continuing advances in computer power, financial data often have a number of additional characteristics. For example, financial data are often considered very 'noisy', which means that it is more difficult to separate *underlying trends or patterns* from random and uninteresting features. Financial data are also almost always not normally distributed in spite of the fact that most techniques in econometrics assume that they are. High frequency data often contain additional ‘patterns’ which are the result of the way that the market works, or the way that prices are recorded. These features need to be considered in the model-building process, even if they are not directly of interest to the researcher.

One of the most rapidly evolving areas of financial application of statistical tools is in the modelling of market microstructure problems. ‘Market microstructure’ may broadly be defined as the process whereby *investors’ preferences and desires are translated into financial market transactions*. It is evident that microstructure effects are important and represent a key difference between financial and other types of data. These effects can potentially impact on many other areas of finance. For example, market rigidities or frictions can imply that current asset prices do not fully reflect future expected cashflows (see the discussion in Chapter 10 of this book). Also, investors are likely to require compensation for holding securities that are illiquid, and therefore embody a risk that they will be difficult to sell owing to the relatively high probability of a lack of willing purchasers at the time of desired sale. Measures such as volume or the time between trades are sometimes used as proxies for market liquidity.

A comprehensive survey of the literature on market microstructure is given by Madhavan (2000). He identifies several aspects of the market microstructure literature, including price formation and price discovery, issues relating to market structure and design, information and disclosure. There are also relevant books by O’Hara (1995), Harris (2002) and Hasbrouck (2007). At the same time, there has been considerable advancement in the sophistication of econometric models applied to microstructure problems. For example, an important innovation was the autoregressive conditional duration (ACD) model attributed to Engle and Russell (1998). An interesting application can be found in Dufour and Engle (2000), who examine the effect of the time between trades on the price-impact of the trade and the speed of price adjustment.

1.3 Steps Involved in Formulating an Econometric Model

Although there are of course many different ways to go about the process of model-building, a logical and valid approach would be to follow the steps described in Figure 1.1.

The steps involved in the model construction process are now listed and described. Further details on each stage are given in subsequent chapters of this book.

- *Steps 1a and 1b: general statement of the problem* This will usually involve the formulation of a theoretical model, or intuition from financial theory that two or more variables should be related to one another in a certain way. The model is unlikely to be able to completely capture every relevant real-world phenomenon, but it should present a sufficiently good approximation that it is useful for the purpose at hand.

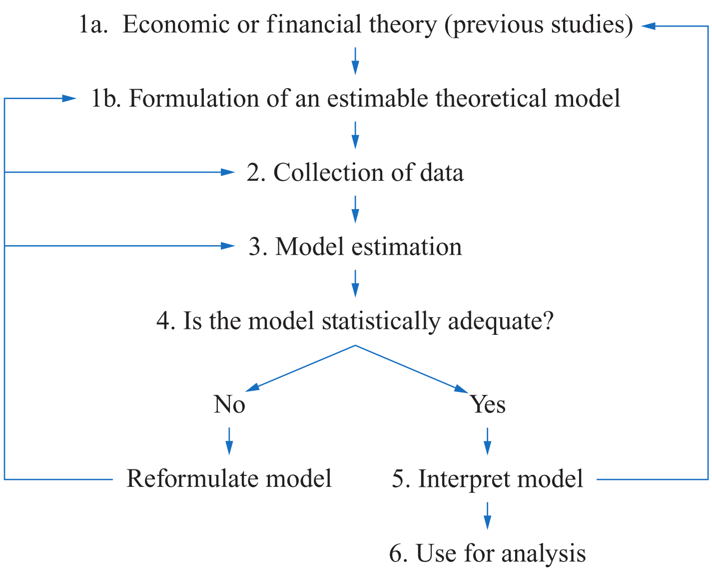

Figure 1.1 Steps involved in formulating an econometric model

- *Step 2: collection of data relevant to the model* The data required may be available electronically through a financial information provider, such as Reuters or from published government figures. Alternatively, the required data may be available only via a survey after distributing a set of questionnaires, i.e., *primary* data.
- *Step 3: choice of estimation method relevant to the model proposed in step 1* For example, is a single equation or multiple equation technique to be used?
- *Step 4: statistical evaluation of the model* What assumptions were required to estimate the parameters of the model optimally? Were these assumptions satisfied by the data or the model? Also, does the model adequately describe the data? If the answer is 'yes', proceed to step 5; if not, go back to steps 1–3 and either reformulate the model, collect more data, or select a different estimation technique that has less stringent requirements.
- *Step 5: evaluation of the model from a theoretical perspective* Are the parameter estimates of the sizes and signs that the theory or intuition from step 1 suggested? If the answer is 'yes', proceed to step 6; if not, again return to stages 1–3.
- *Step 6: use of the model* When a researcher is finally satisfied with the model, it can then be used for testing the theory specified in step 1, or for formulating forecasts or suggested courses of action. This suggested course of action might be for an individual (e.g., 'if inflation and GDP rise, buy stocks in sector X'), or as an input to government policy (e.g., 'when equity markets fall, program trading causes excessive volatility and so should be banned').

Cambridge University Press
978-1-108-42253-6 — Introductory Econometrics for Finance
Chris Brooks
Excerpt
More Information

It is important to note that the process of building a robust empirical model is an iterative one, and it is certainly not an exact science. Often, the final preferred model could be very different from the one originally proposed, and need not be unique in the sense that another researcher with the same data and the same initial theory could arrive at a different final specification.

1.4 Points to Consider When Reading Articles in Empirical Finance

As stated above, one of the defining features of this book relative to others in the area is in its use of published academic research as examples of the use of the various techniques. The papers examined have been chosen for a number of reasons. Above all, they represent (in this author's opinion) a clear and specific application in finance of the techniques covered in this book. They were also required to be published in a peer-reviewed journal, and hence to be widely available.

When I was a student, I used to think that research was a very pure science. Now, having had first-hand experience of research that academics and practitioners do, I know that this is not the case. Researchers often cut corners. They have a tendency to exaggerate the strength of their results, and the importance of their conclusions. They also have a tendency not to bother with tests of the adequacy of their models, and to gloss over or omit altogether any results that do not conform to the point that they wish to make. Therefore, when examining papers from the academic finance literature, it is important to cast a very critical eye over the research – rather like a referee who has been asked to comment on the suitability of a study for a scholarly

BOX 1.2 Points to consider when reading a published paper

(1) Does the paper involve the development of a theoretical model or is it merely a technique looking for an application so that the motivation for the whole exercise is poor?

(2) Are the data of 'good quality'? Are they from a reliable source? Is the size of the sample sufficiently large for the model estimation task at hand?

(3) Have the techniques been validly applied? Have tests been conducted for possible violations of any assumptions made in the estimation of the model?

(4) Have the results been interpreted sensibly? Is the strength of the results exaggerated? Do the results actually obtained relate to the questions posed by the author(s)? Can the results be replicated by other researchers?

(5) Are the conclusions drawn appropriate given the results, or has the importance of the results of the paper been overstated?

journal. The questions that are always worth asking oneself when reading a paper are outlined in Box 1.2.

Bear these questions in mind when reading my summaries of the articles used as examples in this book and, if at all possible, seek out and read the entire articles for yourself.

This chapter now moves on to cover the fundamental mathematical framework that underpins financial econometrics. This material is intended as a refresher for readers who have covered these topics in the past but require a reminder; students who are seeing these concepts for the first time may find a more thorough treatment covering an entire book useful in addition to this text – see, for example Renshaw (2016) or Swift and Piff (2014), which are both detailed and very accessible.

1.5 Functions

1.5.1 Introduction to Functions

The ultimate objective of econometrics is usually to build a model, which may be thought of as a simplified version of the true relationship between two or more variables that can be described by a *function*. A function is simply a mapping or relationship between an input or set of inputs and an output. We usually write that y, the output, is a function f of x, the input, so $y = f(x)$. $f(.)$ is simply a general method of stating that y is related to x in some fashion. Another way to say this is that f provides a mapping between y and x so that it tells us, for every given value of x, what the corresponding value of y would be. f is a unique (1:1) mapping so that for each value of x there is only one corresponding value of y.

The *domain* of x is defined as the set of values that this variable can take; the *range* refers to the respective set of values that y can take. Usually, neither the domain nor the range are specified, in which case they can both be assumed to be allowed to take any real values.

1.5.2 Straight Lines

y could be a linear function of x, where the relationship can be expressed as a straight line on a graph, or y could be a non-linear function of x, in which case the relationship between the two variables would be represented graphically as a curve. If the relationship is linear, we could write the equation for this straight line as

$$y = a + bx \tag{1.1}$$

y and x are called *variables*, while a and b are *parameters*; a is termed the *intercept* and b is the *slope* or *gradient* of the line. The intercept is the point at which the line crosses the y-axis, while the slope measures the steepness of the line. Note that there will be only one value of a and one value of b, although there will be many values of x and of y. a and b could each be any combination of positive, negative or zero.

Table 1.1 Sample data on hours of study and grades

Hours of study (x)	Grade-point average in % (y)
0	25
100	30
400	45
800	65
1000	75
1200	85

To illustrate, suppose we were trying to model the relationship between a student's grade-point average y (expressed as a percentage), and the number of hours that they studied throughout the year, x. Suppose further that the relationship can be written as a linear function with $y = 25 + 0.05x$.

Clearly it is unrealistic to assume that the link between grades and hours of study follows a straight line, but let us keep this assumption for now. So the intercept of the line, a, is 25, and the slope, b, is 0.05. What does this equation mean? It means that a student spending no time studying at all ($x = 0$) could expect to earn a 25% average grade, and for every hour of study time, their average grade should improve by 0.05% – in other words, an extra 100 hours of study through the year would lead to a 5% increase in the grade.

Suppose that a particular student wished to score a perfect 100% grade-point average. How many hours would (s)he need to study? To calculate this, we would need to set $y = 100$ and then to solve for x: $100 = 25 + 0.05x$, so $x = 1500$ hours. We could construct a table with several values of x and the corresponding value of y as in Table 1.1 and then plot them onto a graph (Figure 1.2).

We can see from the graph that the gradient of this line is positive (i.e., it slopes upwards from left to right). Note that for a straight line, the slope is the same along the whole line; this slope can be calculated from a graph by taking any two points on the line and dividing the change in the value of y by the change in the value of x between the two points.

In general, a capital delta, Δ, is used to denote a change in a variable. For example, suppose that we want to take the two points $x = 100$, $y = 30$ and $x = 1000$, $y = 75$. We could write these two points using a coordinate notation (x,y) and so (100,30) and (1000,75) in this example. We would calculate the slope of the line as

$$\frac{\Delta y}{\Delta x} = \frac{75 - 30}{1000 - 100} = 0.05 \tag{1.2}$$

So indeed, we have confirmed that the slope is 0.05 (although in this case we knew that from the start). Two other examples of straight line graphs are given in Figure 1.3.

Cambridge University Press
978-1-108-42253-6 — Introductory Econometrics for Finance
Chris Brooks
Excerpt
More Information

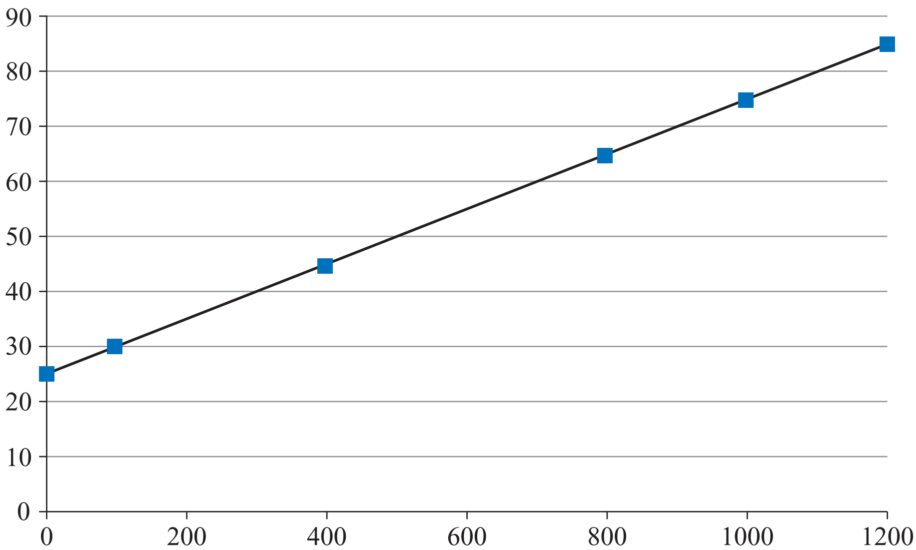

Figure 1.2 A plot of hours studied (x) against grade-point average (y)

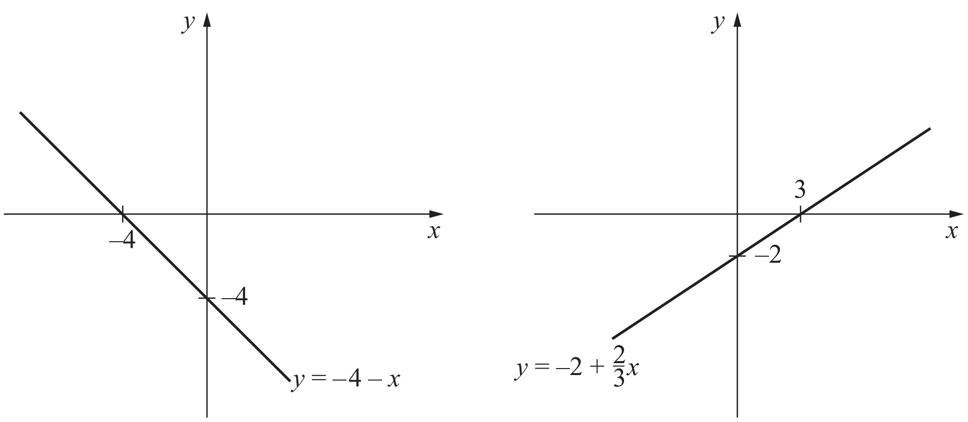

Figure 1.3 Examples of different straight line graphs

The gradient of the line can be zero or negative instead of positive. If the gradient is zero, the resulting plot will be a flat (horizontal) straight line. We could then write it as $y = 25 + 0x$, so that whatever the value of x, y will always be the same (25).

If there is a specific change in x, Δx, and we want to calculate the corresponding change in y, we would simply multiply the change in x by the slope, so $\Delta y = b\Delta x$.

As a final point, note that we stated above that the point at which a function crosses the y-axis is termed the intercept. The point at which the function crosses the x-axis is called its *root*. In the example above, if we take the function $y = 25 + 0.05x$, set y to zero and rearrange the equation, we would find that the root would be $x = -500$.

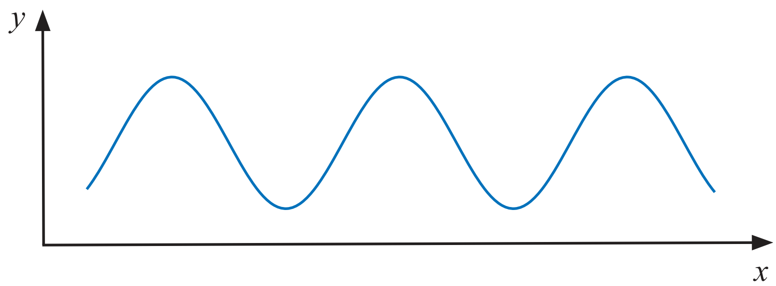

Figure 1.4 Example of a general polynomial function

In this case, the root of the equation does not have a useful interpretation (as the number of hours studied cannot be negative) but this will not always be the case.

The equation for a straight line has one root (except for a horizontal straight line such as $y = 4$, where there would be no root since it never crosses the x-axis). Further examples of how to calculate the roots of an equation will be given in Section 1.5.3.

1.5.3 Polynomial Functions

A linear function is often not sufficiently flexible to be able to accurately describe the relationship between two variables, and so a quadratic function may be used instead. A *polynomial* simply adds higher order powers of the variable x into the function. In the most general case, we would have an n^{th} order polynomial (a polynomial of order n)

$$y = a + b_1 x + b_2 x^2 + b_3 x^3 + \cdots + b_n x^n \tag{1.3}$$

If $n = 2$, we have a quadratic equation, if $n = 3$ a cubic, if $n = 4$ a quartic and so on. We use polynomials if y depends only on one variable x but in a non-linear way (and so it cannot be expressed as a straight line). An example of the shape of a general polynomial function is given in Figure 1.4.

Broadly, the higher the order of the polynomial, the more complex will be the relationship between y and x and the more twists and turns there will be in the plot like Figure 1.4. However, usually $n = 2$, a quadratic equation, is sufficient to describe the function as it seems unlikely that a real series y will rise with x then fall before rising again and so on, which would be the case if it was described by a higher order polynomial. So now we will focus on the quadratic case.

We could write the general expression for a quadratic function as

$$y = a + bx + cx^2 \tag{1.4}$$

where x and y are again the variables and a, b, c are the parameters that describe the shape of the function. Note that we have changed notation slightly for simplicity between equations (1.3) and (1.4), writing the slope parameters as b and c rather than