Preface

Analytic Methods for Managerial Decision Making

This preface helps orient you in the "frame of mind" that lies behind much of this book. Some of its content will become more meaningful after you see applications. Therefore, it is recommended that you read it once before continuing (without attempting to absorb it all) and then return to it later.

P.1 Motives and objectives

Broadly

You manage a firm. The operations of the firm are very complex, with many decisions to be made about production, marketing, financing, and so on. Call a complete configuration of these decisions a *strategy*. Each strategy results in a profit for your firm. Your problem is to choose the strategy with the highest profit.

Here is how to solve this problem. Open up a spreadsheet with two columns. In one column, list the possible strategies. In the other column, list the profit for each strategy. Either by eyeballing the spreadsheet or using one of the spreadsheet program's built-in functions, pick out the highest profit level. In the same row, look at the entry in the strategy column. This is your best strategy.

If only life were so simple! In practice: (a) the problem is very complex—for example, you could not even list all the strategies; and (b) you do not have the hard data needed to determine the profit of each strategy. (Luckily for you—this is why firms hire human managers rather than computers to make decisions.)

The goal of this book is to enhance your ability to make decisions *on your feet* using *soft data* in *complex situations*. In pursuit of this goal, we introduce several methods, central to which are *logical analysis* and *simplification*.

1. Logical analysis is one of the important tools a manager brings to bear on a problem (others include intuition, experience, and knowledge of similar cases).
2. Good managers are smart but do not have infinite information-processing ability. Therefore, like good physicists, good doctors, and good economists, they simplify problems in order to reason logically about them.

More specifically

We cover the following methods, which are used in every topic in this book.

Models. A model is a simplified or stylized description of a problem that isolates the most important features.

Smooth functions. Even when the underlying data is discrete, we may use smooth approximations to simplify our analysis.

Decomposition of decision problems. Decomposing a decision problem means to divide it into smaller and simpler subproblems.

Marginal analysis. Marginal analysis considers the incremental effects of small changes in decisions. Under the right conditions, marginal analysis provides a simple way to find an optimal decision or to check whether a decision is optimal.

P.2 The economist's notion of models

> *Although this may seem like a paradox, all exact science is dominated by the idea of approximation.* — Bertrand Russell

Models

A model is an artificial situation (a "metaphor") that is related to, but simpler than, the real-world situations that the model is meant to help us understand. Every topic in this book is studied by constructing models.

Simplification is a goal of modeling, not an unintended negative consequence. This simplicity has two roles.

1. The human brain cannot comprehend all aspects of a real-world situation simultaneously; hence, it is useful to decompose it into various simpler situations. Only then can we apply logical analysis.
2. Our goal is not to derive conclusions from a wealth of data about a few cases. Instead, we wish to say as much as possible using as little information as possible. This will make it more likely that our conclusions apply to a broad range of future experiences in which you will often have limited information.

So you should judge a model by what is in it, not by what has been left out. The components we ignore in a model might introduce new relationships, but they will not invalidate those we have identified. Of course, all conclusions drawn from models must be taken with a grain of salt rather than applied dogmatically.

In Chapter 1B, for example, we construct a model of a simple market. The model does not include any details about how traders interact and settle on trades. Ignoring such details not only makes our model simpler, it also allows us to draw *robust* conclusions that are relevant to a wide variety of trading mechanisms.

We next discuss three of the simplifying assumptions that appear in most models in this book (though they are not part of all economic models): rationality, no uncertainty, and partial equilibrium. Simplifying assumptions are not "correct"—otherwise they would not be simplifying. Therefore, to discuss such assumptions means to explain what is lost and what is gained by making them.

Rationality

Microeconomics, like the other disciplines that underlie your business education, is a social science. We are interested in understanding the interaction between people. For this, we need to specify how each individual behaves.

People cannot handle unlimited amounts of information and they make mistakes, but that does not mean that their behavior is arbitrary or irrational. In important economic interactions, people are goal oriented and work very hard to pursue these goals. Not all entrepreneurs are equally good at managing a company, but they all work very hard at earning a profit and it is this goal-oriented behavior that is the largest determinant of their actions.

The simplest way to capture such goal-oriented behavior is to ignore the imperfections and limitations in people's abilities. Economists call this the rationality assumption, though "rationality" has a stronger meaning here than in everyday discourse because it implies that each person is infinitely smart and makes no mistakes. (Economists use the term "bounded rationality" for behavior that is goal-oriented and reasonable but with the usual human limitations on processing information.)

Besides being a simple way to capture real-life goal-oriented behavior, the rationality assumption also helps us learn how we should act in economic situations. To make good decisions, we should learn how a hypothetical "rational" person would behave.

No uncertainty

In most of this book, we assume that people know the information relevant to their decision problems. Models with uncertainty are more complex and should only be studied after understanding the corresponding model without uncertainty.

Partial equilibrium

We study *partial equilibrium* models. This means that we concentrate on one market at a time while keeping other activities and prices fixed. Economics also has *general equilibrium* models, which consider the simultaneous determination of all prices. These are more complex and are not well suited to the study of pricing by firms with market power or of strategic interaction between market players.

P.3 Smooth approximations

Most economic variables are not perfectly divisible. For example, prices must usually be quoted in a smallest currency unit and most goods come in multiples of a smallest quantity (fax machines, bars of soap, sheets of paper). However, we can safely ignore these indivisibilities if the smallest unit is small compared to the relevant scale of prices or output. We thus obtain a smooth relationship between economic variables that simplifies much analysis.

Consider Figure P.1 (on page xxiii), which shows a firm's profit as a function of the level of output. The quantity is denoted by Q, the profit by Π, and the profit function (which relates each quantity to its profit) by $\pi(Q)$. (We denote profit by the Greek letter "pi": lowercase π and uppercase Π.) You can see that the optimal output level is only 8 units. Perhaps this firm is a small builder of single-family homes (and profit is measured in \$1000s).

Suppose instead that the firm builds in-ground swimming pools. The size of the unit is smaller compared to the optimal scale of output, so the graph of the profit function could look like Figure P.2.

Once the scale of output is up to, say, 200 or 300 units, which is the case for most firms, this graph will look almost like a smooth line. It is then a reasonable approximation to treat the good as perfectly divisible (like oil or water), so that the graph looks like Figure P.3.

P.4 Decomposition of decision problems

Decomposing a decision problem means dividing it into several smaller and simpler problems. This is useful because it is easier to understand one simple problem at a time than an entire complex problem all at once. Furthermore, one can distribute the task of making decisions among the many members of an organization, thereby making use of more diversified brain power.

For example, suppose your task is to choose the output level Q of your single-product firm. As a simple accounting identity, profit equals revenue minus cost: $\Pi = R - C$. To express this as a function of Q, we let $r(Q)$ be the highest possible revenue when selling Q units and let $c(Q)$ be the total cost when the output level is Q. Then $\pi(Q) = r(Q) - c(Q)$.

We can see how maximizing profit involves trading off revenue and cost. Increasing output from Q_1 to Q_2 raises profit if and only if the increase in revenue exceeds the increase in cost: $r(Q_2) - r(Q_1) > c(Q_2) - c(Q_1)$.

Figure P.1

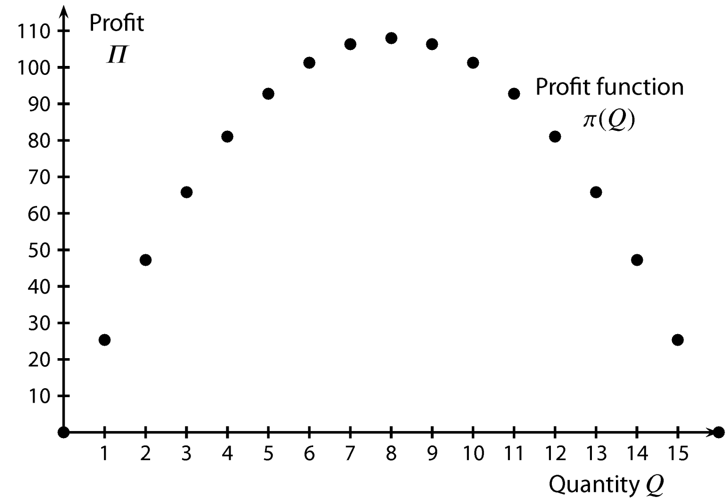

Figure P.2

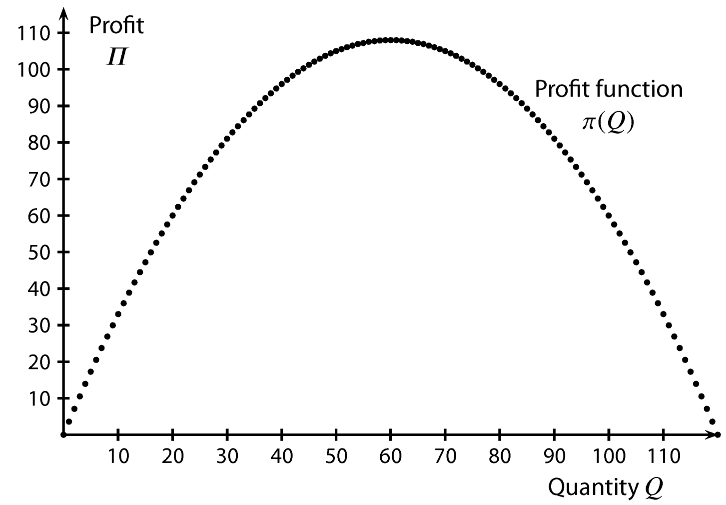

Figure P.3

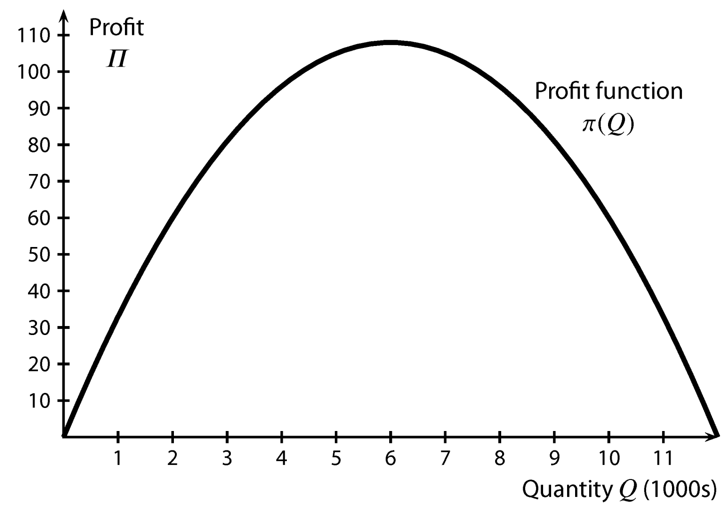

Furthermore, we can delegate the task of determining revenue to the marketing department and we can delegate the task of determining cost to the production department. For example, the marketing department must determine how to market the output and what price to charge in order to achieve the highest revenue when selling Q units. We can then set the output level knowing only the functions $r(Q)$ and $c(Q)$, without having to know the details of the tasks we delegated to the marketing and production departments.

P.5 Marginal analysis

Firms are not frictionless reflections of their momentary environments, but rather highly inertial action repertoires, responding to—indeed perceiving—today's environment largely in terms of lessons learned from actions in days gone by.

Workshop at the Santa Fe Institute[1]

As with any new technology, the early years of the Internet have been a learning process—and here's what we now know. First, the Internet was supposed to change everything. That's just plain wrong. Clearly, in much of the economy, the Internet offers incremental payoffs without substantially altering core businesses.

Article in *Business Week Online*[2]

Suppose you come into an organization as a consultant or as a new CEO. Your job is to make this company perform as well as possible. You should start from scratch, right? Don't even look at the current state of affairs. That's what the "change gurus" say: "The larger the scale of change, the greater the opportunity for success" (James Champy).

This sounds exciting, but it ignores the complexity of large-scale change. Shying away from complexity is not a sign of weakness. It is a sign of wisdom, of recognizing an iron law: The more complex a problem is, the more likely mistakes are and the more costly it is to avoid them. This section is about how to make use of the simplicity of incremental change and how to recognize when it works.

1. Michael Cohen, Roger Burkhart, Giovanni Dosi, Massimo Egidi, Luigi Marengo, Massimo Warglien, and Sidney Winter, "Routines and Other Recurring Action Patterns of Organizations: Contemporary Research Issues", *Industrial and Corporate Change*, 1996.

2. Michael Mandela and Robert Hof, "Rethinking the Internet", *Business Week Online*, 26 March 2001.

Marginal conditions: an allegory

Suppose a blind man wants to get to a peak on an island, such as the one drawn in Figure P.4.

Figure P.4

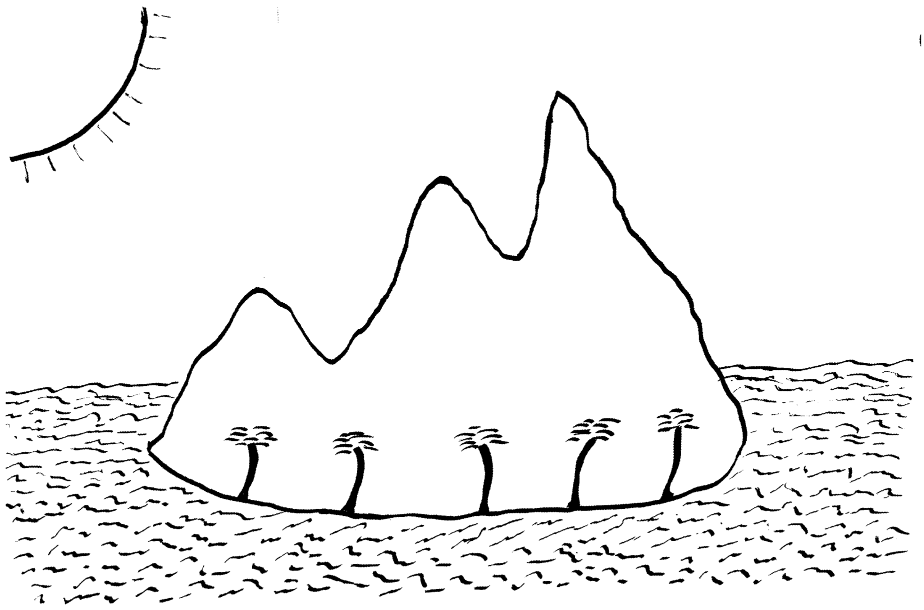

To *check* whether he is at a peak, he need only tap his cane to see if any point around him is higher—if not, then he is at the top of a peak. This is called the *local* or *marginal condition for optimality*. To *reach* a peak from any point on the island, he can tap his cane to see in which direction the slope goes up and then take his next step in that direction, repeating until there is nowhere higher to go. This is called *local* or *incremental search*.

Such local methods are useful to the blind man. However, on the island in Figure P.4, the blind man cannot be sure to find the highest possible peak. On the other hand, local methods work perfectly on the island in Figure P.5 because it has a single peak, with the rest of the island sloping up toward it.

Figure P.5

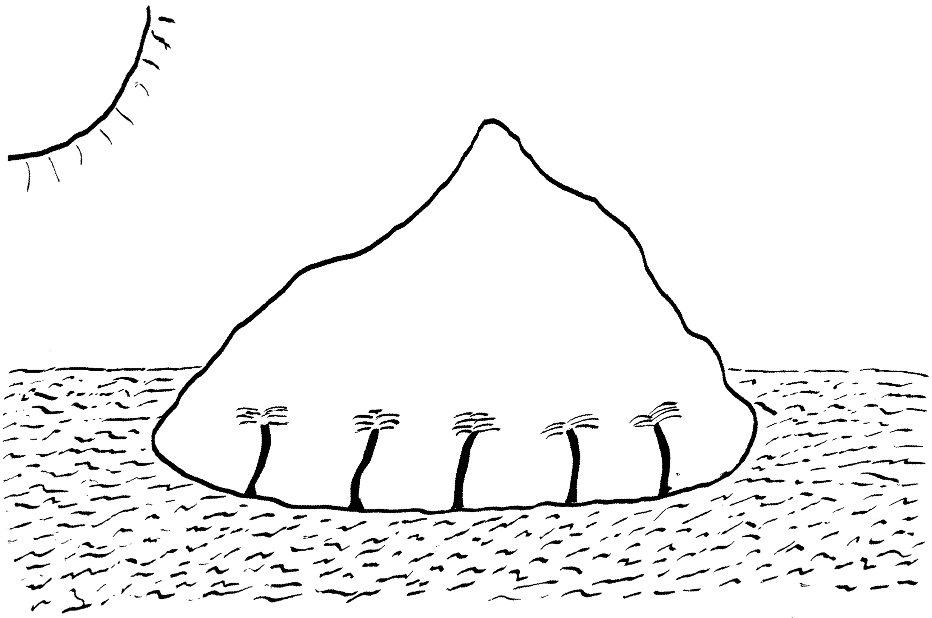

So, although local methods are useful on both islands, only on the second island can they be used on their own to find the highest point on the island. Therefore, we say that *marginal conditions are sufficient* on the second island, but not on the first.

Marginal conditions: a manager

Let's relate the story about the blind man to managerial decision making. Suppose you manage a firm that produces a single indivisible good and you have to decide how much to produce and sell in order to maximize the firm's profit $\pi(Q)$. Profit is zero if you produce and sell nothing (zero cost and zero revenue). The firm makes a loss if you produce too much because, in order to sell the output, you have to lower your price to below the average cost of production. For intermediate output levels, you can make a profit. The question is: how much should you produce?

Suppose that, for the first 16 units, the graph of the profit function is as shown in

Figure P.6

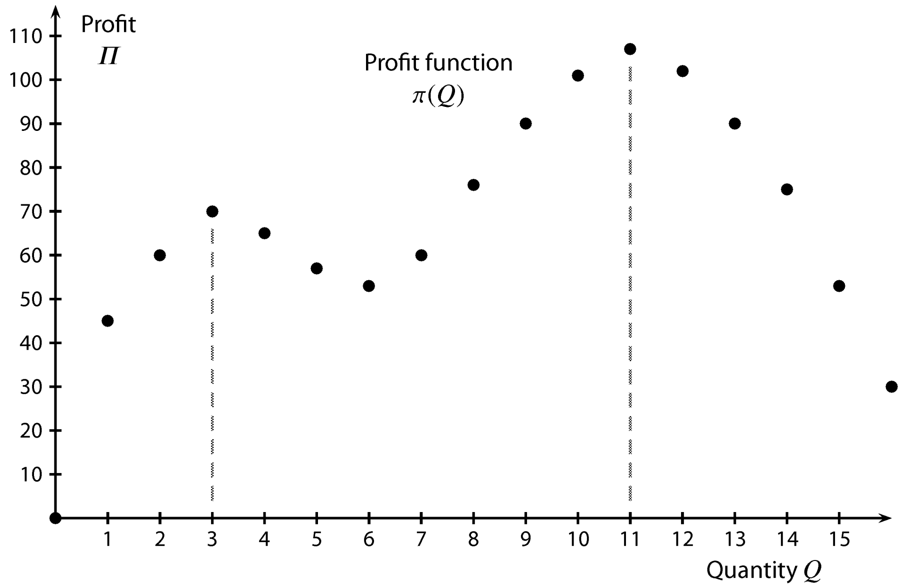

The two peaks, at $Q = 3$ and $Q = 11$, are called *local optima* because each gives a higher profit than nearby quantities. The quantity $Q = 11$ is the *global optimum*—it gives you the overall highest profit.

Realistically, you do not actually know the exact relationship between quantity and profit. You can determine the profit for particular quantities either by experimenting or by acquiring information and performing calculations, but doing so for all quantities would be costly. Thus, you are like the blind man: you can feel the terrain around you but you do not know the topography of the entire island.

Still, you can at least reach a local optimum using marginal conditions and local search.

Marginal conditions. To check whether a level Q is a local optimum, it suffices to check that the profit is not higher at the *nearby* alternatives $Q - 1$ and $Q + 1$. Only at $Q = 3$ and $Q = 11$ will you find that neither of the two nearby quantities yields a higher profit.

Local search. To find a locally optimal level of sales, you can search *nearby* for improvements until you find no better alternatives. Suppose you start with output equal to 5. You check the quantities 4 and 6 and determine that 4 gives you a higher profit. Then you check 3 and find it is also better. Continuing, you find that decreasing your output to 2 does not lead to a further improvement. Hence, you have found that 3 is a local optimum.

If your profit function is as shown in Figure P.6, such local or marginal analysis is useful for finding and identifying local optima, but it does not guarantee, by itself, that you find the profit-maximizing output level (the global optimum). In the preceding example of local search, you set the output to 3, unaware that a jump in your output to 11 would give a higher profit.

However, marginal analysis works perfectly if your profit function has a single peak, as in Figure P.7. Then every local optimum is a global optimum; we say that *marginal conditions are sufficient*.

Figure P.7

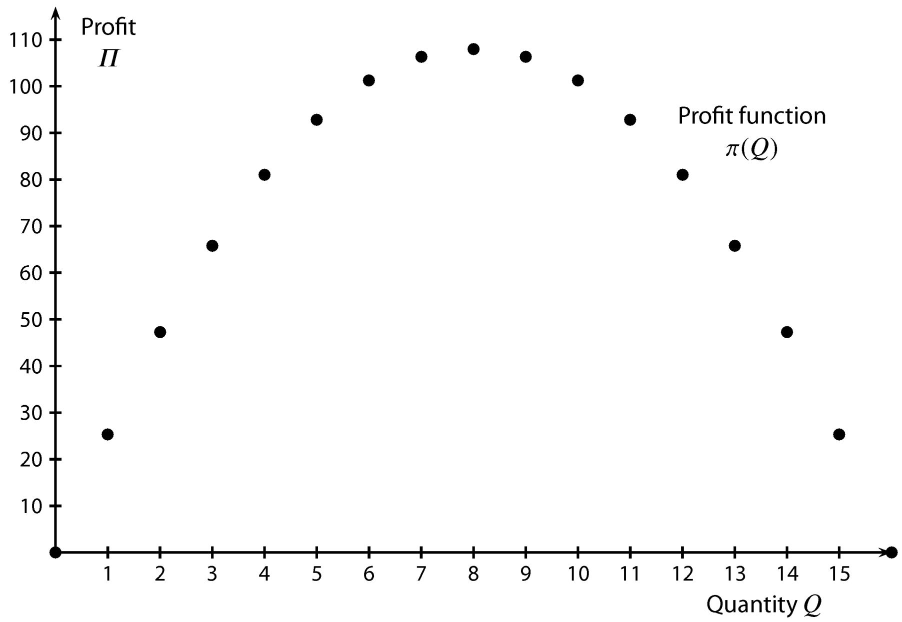

In this book, we maintain an unstated assumption that marginal conditions are sufficient, except in a few cases in which we state otherwise: a competitive firm with a U-shaped average cost curve, a firm with market power that has a fixed cost; and a consumer that faces a two-part tariff. In these cases, the only required supplement to marginal analysis is that you have to check a shut-down option.

P.6 The mathematics of marginal analysis

Marginal profit

Consider again the problem of chooses output in order to maximize profit $\pi(Q)$. "Marginal conditions" mean conditions that must be satisfied in order for a quantity Q to be a local optimum. We state marginal conditions in terms of the marginal profit, which is defined differently depending on whether we have a discrete or smooth function.

$$\text{Discrete:} \quad m\pi(Q) = \pi(Q) - \pi(Q-1).$$
$$\text{Smooth:} \quad m\pi(Q) = \pi'(Q).$$

Marginal conditions in the discrete case

Consider the marginal conditions in the discrete case. In Figure P.6, $\pi(2) = 60$, $\pi(3) = 70$, and $\pi(4) = 65$. Therefore,

$$m\pi(3) = \pi(3) - \pi(2) = 70 - 60 = 10,$$
$$m\pi(4) = \pi(4) - \pi(3) = 65 - 70 = -5.$$

The fact that $m\pi(3) \geq 0$ tells us that lowering output from 3 to 2 would cause profit to fall (or stay the same); the fact that $m\pi(4) \leq 0$ tells us that raising output from 3 to 4 would cause profit to fall (or stay the same). Thus, together these conditions tell us that $Q = 3$ is at least a local optimum.

In summary, the marginal conditions for Q to be a local optimum are as follows:

$$m\pi(Q) \geq 0 \quad \text{and} \quad m\pi(Q+1) \leq 0.$$

Thus, a local optimum is a point where marginal profit shifts from positive to negative. In Figure P.7, this happens at $Q = 8$; in Figure P.6, this happens at $Q = 3$ and $Q = 11$.

Marginal conditions in the smooth case

When the curve is smooth, as in Figure P.8, the slope of the profit function is exactly zero at a local optimum; this is illustrated by the horizontal line drawn in Figure P.8.

Figure P.8

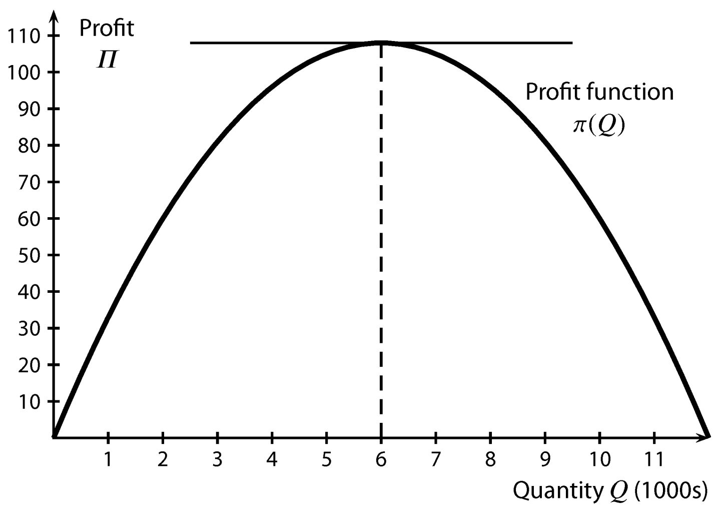

That is, the marginal condition for an optimum is $m\pi(Q) = 0$. This is simpler than the pair of inequalities $m\pi(Q) \geq 0$ and $m\pi(Q+1) \leq 0$ in the discrete case, which is one reason that smooth functions are easier to work with.

Using marginal conditions for a numerical example

Though the real purpose of the marginal analysis (and our use of calculus for marginal analysis) in this text is to draw qualitative conclusions, marginal conditions can also be used to solve numerical examples. Suppose that a firm has the following revenue and cost curves:

$$r(Q) = 60Q - 3Q^2,$$
$$c(Q) = 24Q.$$

Then the profit curve is

$$\pi(Q) = r(Q) - c(Q) = (60Q - 3Q^2) - 24Q = 36Q - 3Q^2.$$

Taking the derivative yields the marginal profit curve:

$$m\pi(Q) = 36 - 6Q.$$

We solve the marginal condition (also called the first-order condition in mathematics):

$$\begin{aligned} m\pi(Q) &= 0 \\ 36 - 6Q &= 0, \\ 36 &= 6Q, \\ Q &= 6. \end{aligned}$$

Thus (assuming that marginal conditions are sufficient, which is true in this case), the profit-maximizing quantity is $Q = 6$. This is the data that lies behind Figure P.8, in which one can see that $Q = 6$ maximizes profit and that $m\pi(6) = 0$.

Decomposition of marginal conditions

When we decompose a decision problem, we can also decompose the marginal conditions. Since $\Pi = R - C$, it follows that $M\Pi = MR - MC$ and that the marginal conditions can be restated as in Table P.1.

Table P.1

	Marginal conditions	
	In terms of marginal profit	In terms of marginal revenue and cost
Discrete case	$m\pi(Q) \geq 0$ $m\pi(Q+1) \leq 0$	$mr(Q) \geq mc(Q)$ $mr(Q+1) \leq mc(Q+1)$
Smooth case	$m\pi(Q) = 0$	$mr(Q) = mc(Q)$

In words, we usually state the marginal conditions as "marginal revenue equals marginal cost". This is exact for the smooth case and approximate for the discrete case.

Let's illustrate such decomposition using data from our previous example:

$$r(Q) = 60Q - 3Q^2,$$
$$c(Q) = 24Q,$$
$$\pi(Q) = 36Q - 3Q^2.$$

Figure P.9 shows these three curves on the same graph.

Figure P.9

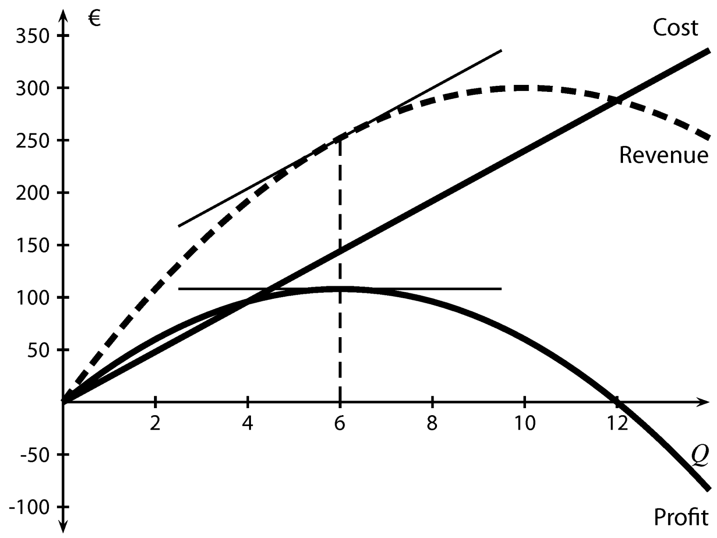

Observe that, at the profit-maximizing quantity $Q = 6$, where $m\pi(Q) = 0$, it is also true that the revenue and cost curves have the same slope—that is, marginal revenue equals marginal cost.

We can solve numerical examples by solving the marginal condition $mr(Q) = mc(Q)$. Taking the derivative of the revenue and cost curves yields

$$mr(Q) = 60 - 6Q,$$
$$mc(Q) = 24.$$

(The cost curve is linear and so has the constant slope 24.) We solve

$$mr(Q) = mc(Q)$$
$$60 - 6Q = 24$$
$$36 = 6Q$$
$$Q = 6.$$

Of course, we get the same answer $Q = 6$ as when we solved $m\pi(Q) = 0$.

Exercise P.1. Suppose a firm's revenue and cost curves are

$$r(Q) = 84Q - 3Q^2,$$
$$c(Q) = 4Q + Q^2.$$

a. Write the formula for the profit curve.

b. Using a spreadsheet, create a table with 4 columns: Output (ranging from 0 to 20), revenue, cost, and profit. By visual inspection of the profit column, determine the quantity Q^* that maximizes profit.

c. Using your spreadsheet program (or another graphing program), graph $r(Q)$, $c(Q)$, and $\pi(Q)$. Print out the graph. Mark the point where profit is maximized. Draw lines tangent to the graphs of $r(Q)$, $c(Q)$, and $\pi(Q)$ at the profit-maximizing Q^*. You should see visually that $m\pi(Q^*) = 0$ and that $mr(Q^*) = mc(Q^*)$.

d. Calculate $m\pi(Q)$ and find the profit-maximizing quantity by solving $m\pi(Q) = 0$.

e. Calculate $mr(Q)$ and $mc(Q)$. Find the profit-maximizing quantity by solving $mr(Q) = mc(Q)$.

P.7 Wrap-up

Our use of models and our analysis of decision problems are based on *simplification*. Only through simplification can we carefully describe the key elements of a situation and then apply logical reasoning.

The key tools for simplifying decision problems are decomposition and marginal analysis. These are important not only for analyzing other people's decisions but also as methods for making decisions ourselves.

Such logical analysis can be used to develop quantitative decision-making tools. However, it is also useful for making decisions on your feet using soft data. Improving such decision making is the ultimate goal of this book.