Chapter 8: A Toolkit: Game Theory and Imperfect Competition Models

February 14, 2003

1 Introduction

This chapter aims at introducing the reader to imperfect competition models that are used in the technical sections of the book (to be more precise, the intermediate technical sections, labelled with *; the advanced technical sections, labelled **, will likely need a stronger background than the one which is offered in this chapter).

Of course, this is no replacement for a more proper training in basic industrial organisation models, but hopefully it will help some students who have already some background in economics and in simple mathematical analysis (I use little more than derivatives of real functions in the book and in this chapter), or those who want to refresh a knowledge acquired some time ago, to follow the formal arguments made in the book.

The choice of the topics analysed here is functional to the objective of helping the reader follow the book: there are several interesting topics that I would teach in an industrial organisation course but that I am not covering here because they have no (or marginal) appearance in the book.

The section starts with a short treatment of monopoly (section 2), then introduces the reader to elementary game theory (section 2.2.3), which is indispensable to understand modern oligopoly theory, that for convenience I divide into static models (section 3) and dynamic models (section 4).

2 Monopoly

This section offers an introductory treatment of monopoly pricing. First, the case of a single-product monopoly, and then that of a multi-product monopoly are analysed.

2.1 Single-product monopoly

The easiest possible model of imperfect competition is one where there is a monopolist that sells only one good. I will first solve the monopolist's problem with general cost and demand functions, and then offer some specific examples.

Denote demand for this good as $q = D(p)$, where p is the price and q output, and assume that demand is negatively sloped: $\partial D/\partial p < 0$.

For later use, define the *elasticity of demand* ε as the percentage change in the quantity demanded by consumers that follows a one percent change in price: $\varepsilon = -(\partial D/D)/(\partial p/p) = -p(\partial D/\partial p)/D > 0$. (Note that the elasticity is defined as a positive number).

Production costs are given by the function $C(q)$, with non-negative marginal costs: $\partial C/\partial q \geq 0$. For the second-order conditions to be satisfied (that is for the profit function to be concave), assume also that $\partial^2 D/(\partial p)^2 \leq 0$ and $\partial^2 C/(\partial q)^2 \geq 0$.[1]

Let us now see what are the optimal price and quantity set by the monopolist in this market. Its objective is to choose the price that maximises its profits, that is $\max_p \pi = pq - C(q)$, or:

$$\max_p \pi = pD(p) - C(D(p)). \tag{1}$$

The first-order condition (FOC) of this problem is:

$$\frac{\partial \pi}{\partial p} = p\frac{\partial D(p)}{\partial p} + D(p) - \frac{dC(q)}{dq}\frac{\partial D(p)}{\partial p} = 0, \tag{2}$$

that can be re-written as:

$$p - \frac{dC(q)}{dq} = -\frac{D(p)}{\partial D(p)/\partial p}. \tag{3}$$

By using the definition of demand elasticity and substituting, one obtains:

$$\frac{p - dC(q)/dq}{p} = \frac{1}{\varepsilon}. \tag{4}$$

The left-hand side (LHS) of condition (4) is the so-called Lerner index, a measure of market power, that is the ability of a firm to set prices above marginal costs.[2] This condition therefore tells us that the higher the elasticity of demand the lower the monopolist's market power (that is, the lower the relative mark-up it earns). Consider for instance the two extreme cases: if market demand was so inelastic that consumers would be willing to buy the good at whatever price ($\varepsilon = 0$), then the monopolist's mark-up would tend to infinity. If instead demand was extremely elastic, so that consumers would switch to some other goods whenever the monopolist tries to increase marginally the price of its product ($\varepsilon \to \infty$), then its market power would be nil, and the monopolist's price would be equal to its marginal cost (since $p - dC(q)/dq$ must equal zero).

[1] These assumptions are satisfied when adopting a linear demand function and constant marginal costs, which is the case I most often use in this chapter and throughout the book.

[2] See chapter 2 for a discussion of market power (of which a monopoly is an extreme expression) and its relationship with welfare. See chapter 3 for a discussion of how to measure market power.

A specific functions example Let me now use explicit demand and cost functions to analyse the monopolist's problem. Consider for the sake of simplicity the case of a linear demand $q = 1 - p$ and of a constant marginal cost $C(q) = cq$, with $c < 1$ to guarantee viability of the market (otherwise, nobody would buy the product even if the monopolist offered it at marginal cost).

The monopolist's problem is $\max_p \pi = (p - c)(1 - p)$. The FOC is given by:

$$\frac{\partial \pi}{\partial p} = 1 - 2p + c = 0, \tag{5}$$

whence:

$$p^M = \frac{1 + c}{2}. \tag{6}$$

It is worth noting that the higher the marginal cost the higher the price the monopolist will set at equilibrium, $dp^M/dc > 0$: for consumers, better to face an efficient monopolist than an inefficient one![3]

By substitution, one can also find optimal quantities and profits: $q^M = (1 - c)/2$, and $\pi^M = (1 - c)^2/4$.

Note also that whether the monopolist chooses price or quantity the market outcome will be the same (a property that does not hold in oligopoly markets, as we shall see below). To see this, write the inverse demand function as $p = 1 - q$. The monopolist's problem is then $\max_q \pi = (1 - q - c)q$. The FOC is $\partial\pi/\partial q = 1 - 2q - c = 0$, that becomes $q^M = (1 - c)/2$ (and by substitution $p^M = (1 + c)/2$).

2.2 Multi-product monopoly

In the real world, firms often produce and sell more than one product. It is natural, therefore, to ask how a multi-product monopolist prices its different products. I start with the simplest case, where demand and cost of one product does not affect demand and costs for other products, and move later to the cases of interdependent demands and, in turn, interdependent costs. For simplicity, and without losing insights, I restrict attention to the case of two products.

2.2.1 Independent demand and cost functions

Suppose the monopolist sells two products, 1 and 2, that affect neither each other's demands nor costs. The monopolist's total profit is given by: $\pi = \pi_1 + \pi_2 = p_1 D_1(p_1) - C_1(D_1(p_1)) + p_2 D_2(p_2) - C_2(D_2(p_2))$. The monopolist's problem is given by $\max_{p_1,p_2} \pi$, but this decomposes in two completely separate problems: $\max_{p_1} \pi_1$ and $\max_{p_2} \pi_2$. Each of these problems reduces to the problem studied above for the single-product monopolist, giving as a result condition (4):

$$\frac{p_i - dC_i(q_i)/dq_i}{p_i} = \frac{1}{\varepsilon_i}, \quad \text{with } i = 1, 2. \tag{7}$$

[3] See Tirole (1988: 66-67) for a general proof of this result.

This condition tells us that the monopolist sets the price in each market i according to its demand elasticity ε_i: in a market characterised by higher demand elasticity, consumers will pay less than in the other. This is a simple application of a principle of price discrimination that is analysed more at length in chapter 7.

Of course, the assumptions that the price for one good does not affect demand for the other good, and that production costs are independent makes this case very particular. Let us now turn to the cases where interdependence exists, first on the demand side, then on the cost side.

2.2.2 Interdependent demands

Very often, a firm sells a range of products that are to some extent substitutable with each other. Think for instance of a car manufacturer that offers cars with different engine powers, colours and versions (station-wagon, sedan, cabriolet), or a food producer that offers different types of pasta. For substitute products, an increase in the price of one product will increase the demand for the others, that become relatively more convenient: $dq_i/dp_j > 0$.

At times, firms might sell instead (or in addition) products that are complements to each other: a telecom operator might sell both a cellular phone and a subscription to its services, and a food producer might sell both pasta and prepared sauces. For complementary goods, an increase in the price of one product will decrease the demand for the others (by decreasing the demand for the first product, demand for complements is also discouraged): $dq_i/dp_j < 0$.

To see how a monopolist sets prices for two products whose demand are interdependent, assume products have the following demand functions:

$$q_i = a - bp_i + gp_j, \quad \text{with } i, j = 1, 2, \; i \neq j. \tag{8}$$

If the parameter $g > 0$, then the two products are substitutes; if $g < 0$, they are complements; if $g = 0$, their demands are independent of each other (a case that we shall use as a useful benchmark). Assume also that $|g| < b$, to guarantee that the own-price effect on demand of a product is stronger than the cross-price effect, a natural assumption; and that $a > c(b - g)$, to ensure that output is positive at equilibrium.

Assume also that the costs of producing one good is not affected by how much is produced of the other, to focus on the interdependence of demands. More particularly, assume that $C(q_1, q_2) = cq_1 + cq_2$.

The monopolist's total profits are given by:

$$\pi = (a - bp_1 + gp_2)(p_1 - c) + (a - bp_2 + gp_1)(p_2 - c); \tag{9}$$

its problem being $\max_{p_1,p_2} \pi$, the FOCs are:

$$\frac{d\pi}{dp_i} = a - 2bp_i + 2gp_j + c(b - g) = 0, \quad \text{with } i, j = 1, 2; \; i \neq j. \tag{10}$$

At the symmetric solution, $p_1 = p_2 = p_m$, we have:

$$p_m = \frac{a + c(b-g)}{2(b-g)}, \tag{11}$$

where $p_m > c$. Since we are interested in the effect of demand relationship on equilibrium price, note that:

$$\frac{\partial p_m}{\partial g} = \frac{a}{2(b-g)^2} > 0. \tag{12}$$

As g increases in the interval $(-b, b)$, the price charged by the monopolist on both products also increases.[4]

Relative to the benchmark case where the two products are independent ($g = 0$), this implies that the monopolist reduces the price of its products when they are complements ($g < 0$) and it increases them when they are substitutes ($g > 0$). The intuition for this result is straightforward. When the products are complements, they exercise a positive externality on each other and the monopolist internalises it by decreasing its prices (a lower price of good 1 stimulates sales of good 2 and vice versa). In other words, if products 1 and 2 were sold by two distinct monopolists, consumers would pay more for them than when they are sold by the same firm, a result that dates back to Cournot (1838), and that is also discussed both in chapters 6 (a vertically integrated firm being a particular case of a firm that sells complements) and 7 (when studying tie-in sales).

When instead the products are substitutes, the externality they exercise on each other is negative, and the monopolist controls for it by raising prices (a lower price of good 1 crowds out sales of good 2 and vice versa). If products 1 and 2 were sold by two distinct firms, consumers would pay less than when they are sold by the same firm.

To complete the analysis, use the equilibrium price (11) and substitute into output and profit function to obtain:

$$q_m = \frac{a - c(b-g)}{2}; \quad \pi_m = \frac{[a - c(b-g)]^2}{2(b-g)}, \tag{13}$$

q_m being the output per-product and π_m being the total profits.

A dynamic interpretation of demand interdependence So far, I have treated the two products as two distinct products sold simultaneously by a monopolist. However, the insights obtained above carry over to the case where the monopolist sells the same product in sequential markets, simply by re-interpreting the demand relationship as one of inter-temporal substitutability or complementarity.

[4] It is easy to check that p_m is concave, with its lowest possible value (equaling $(a + 2bc)/(4b)$) obtained as g tends to $-b$ and an asymptote as g tends to b.

Continue to consider the case where each good is produced at constant marginal cost c, but assume that demand in the first period is given by $q_1 = a - bp_1$ and demand in the second period by $q_2 = a - bp_2 + \lambda q_1$. If $\lambda > 0$, then higher sales in the first period stimulate demand in the second; if $\lambda < 0$, the opposite holds. The monopolist's total profits (assuming for simplicity a zero interest rate, that is the future gains count as the current ones) are:

$$\pi = (a - bp_1)(p_1 - c) + (a - bp_2 + \lambda(a - bp_1))(p_2 - c). \tag{14}$$

The monopolist chooses p_1 and p_2 so as to maximise π, therefore the two FOCs are given by $d\pi/dp_i = 0$. Solving the resulting system gives:

$$p_1 = \frac{a(1-\lambda) + cb}{b(2-\lambda)}, \quad p_2 = \frac{a + cb(1-\lambda)}{b(2-\lambda)}. \tag{15}$$

It is easy to see that as λ rises the first period price goes down and the second period price goes up:

$$\frac{\partial p_1}{\partial \lambda} = -\frac{a - cb}{b(2-\lambda)^2} < 0, \quad \frac{\partial p_2}{\partial \lambda} = \frac{a - cb}{b(2-\lambda)^2} > 0. \tag{16}$$

The effect of the inter-temporal externality over prices is different. The second period price varies with λ only because it shifts up (or down, if $\lambda < 0$) its demand intercept from a to $a + \lambda q_1$. More interesting is the reason why the first period price varies.

Intertemporal complementarity: Introductory price offers When $\lambda > 0$, the monopolist realises that there is a positive intertemporal demand externality, and it internalises it by decreasing the price relative to the price it would set if there was no future market. In other words, it anticipates that by lowering the first period price it would increase first period demand, that in turn will stimulate demand in the second period. This is an example of a well-known business strategy, that consists of *promotional pricing*, or introductory pricing: in the early stages of the life of a product, the firm sets initially a low price that is progressively increased as consumers get to know and appreciate the product (goodwill effect) or as more consumers have already bought the product (network effects, see chapter 2 for a discussion).[5,6]

Intertemporal substitutability: Durable goods In the case where $\lambda < 0$, higher sales in the first period decrease the demand in the second period.

[5] Another reason why a firm might offer introductory prices is the existence of switching costs (see chapter 2), but while the simple model used here can be seen as a reduced form of a more sophisticated model containing some goodwill effects or network effects, it cannot be easily interpreted as the result of switching costs.

[6] In the particular example chosen here, the first period price does not go below marginal costs. However, one could find examples where the firm charges below marginal cost in the first period.

As a result, the monopolist keeps the first period price higher than in the hypothetical case where the product is sold only once, to internalise the negative demand externality arising across periods. This situation can be seen as a reduced form of the case of a durable good monopolist: the more consumers buy in the first period the fewer will buy in the second. Note that it can be shown that in a durable good monopoly equilibrium prices tend to decrease over time, as is the case in this example.[7]

2.2.3 Interdependent costs: Economies (or diseconomies) of scope

Let us now turn to the impact that cost externalities have over pricing of a multi-product monopolist. Assume that the overall cost function of the monopolist is given by $C(q_1, q_2) = cq_1 + cq_2 + \mu q_1 q_2$.[8] When $\mu > 0$, there exist diseconomies of scope between the two products, as the higher the output of one product the higher the marginal cost of the other product ($\partial^2 C/\partial q_i \partial q_j = \mu > 0$). This is the case, for instance, where both products make use of limited natural resources or inputs: increase in output of one product exerts pressure on the common input by driving up its cost.

When $\mu < 0$, there exist instead economies of scope (or of joint production) between the two products: the higher the output of one product the lower the marginal cost of the other ($\partial^2 C/\partial q_i \partial q_j = \mu < 0$). There are many instances in the real world where producing two goods jointly gives rise to cost savings relative to the case where each product is produced separately.

For simplicity, assume that products are independent, so that $q_i = a - bp_i$.

The monopolist's total profits are:

$$\pi = (a - bp_1)p_1 + (a - bp_2)p_2 - cq_1 - cq_2 - \mu q_1 q_2. \tag{17}$$

Its problem amounts to $\max_{p_1,p_2} \pi$, so the FOCs are:

$$\frac{d\pi}{dp_i} = a(1 + b\mu) - 2bp_i - b^2 \mu p_j + cb = 0, \quad \text{with } i, j = 1, 2;\ i \neq j. \tag{18}$$

Under symmetry, $p_1 = p_2 = p_m$, and the solution becomes:

$$p_m = \frac{a(1 + b\mu) + cb}{b(2 + b\mu)}. \tag{19}$$

To see the effect of the cost externality on the equilibrium price, write:

$$\frac{\partial p_m}{\partial \mu} = \frac{a - bc}{(2 + b\mu)^2} > 0. \tag{20}$$

[7] See chapter 2 for a brief discussion of the durable good monopoly, and chapter 6 for an application that shares many features with the durable good monopoly case (the firm is hurt by its inability to commit to a certain price).

[8] Assume also that $\mu > -2c/a$. This guarantees that marginal costs are positive when μ is negative.

The stronger the cost externality between the two products, the higher the equilibrium price set by the monopolist. In particular, note that the function $p_m(\mu)$ is increasing over all its domain. This implies that when there exist economies of scope ($\mu < 0$), prices are lower than in the benchmark case where the production costs are independent. The firm anticipates that there exists a positive externality between the two products, and reduces the price of each good in order to stimulate its output, and in turn the output of the other product, through the cost reduction. In this case, a multi-product monopolist charges lower prices than two distinct monopolists, each producing one product, would charge.

When instead there exist diseconomies of scope ($\mu > 0$), prices are higher than in the case where production costs are independent. Here there exists a negative externality between the two products, and the monopolist increases the price of each good to internalise it (by reducing the output of one product the marginal cost of the other is lowered). In this case, a multi-product monopolist charges higher prices than two distinct monopolists would charge.

An intertemporal example: learning-by-doing The cost of production of many goods and services decreases with the experience accumulated in producing those goods and services. This is a phenomenon known as "learning-by-doing" (see Arrow, 1961), and that provides an example of inter-temporal externalities on the cost side. As one can see from the following simple example, a monopolist that expects a learning effect will want to decrease prices (relative to a situation where such learning effects are absent) in the early stages of life of a product, in order to increase output and "go down the learning curve", which in turn will make it more efficient in later periods.

Suppose that in each of the two periods of life of its product a monopolist faces demand $p_t = 1 - q_t$, with $t = 1, 2$. Because of learning effects, marginal costs decrease with past production: they are $C'_1 = c$ in the first period and $C'_2(q_1) = c - lq_1$ in the second.[9]

Total profits of the monopolist (assuming there is no discounting for simplicity) are

$$\pi = (p_1 - c)q_1 + (p_2 - c + lq_1)q_2. \tag{21}$$

Since it has to choose q_1, q_2 so as to maximise π, the FOCs are given by $d\pi/dq_1 = 0$ and $d\pi/dq_2 = 0$, that can be re-written as:

$$q_1 = \frac{1 - C'_1 + lq_2}{2}; \quad q_2 = \frac{1 - C'_2}{2}. \tag{22}$$

The previous expressions tell us that while in the second period (the last one in this simple example) the monopolist will behave as usual, that is by simply equating marginal revenue to marginal cost, in the first period it will produce

[9] Assume also that l is positive but small enough, to guarantee that costs are positive in the second period.

more than it would in a static setting, since it internalises the positive externality that this will have on its second-period costs (and the higher l the higher the first period quantity).

By solving the system of FOCs (22) one obtains equilibrium values as:

$$q_1^* = q_2^* = \frac{1-c}{2-l}; \quad p_1^* = p_2^* = \frac{1+c-l}{2-l}. \tag{23}$$

Note that in this simple example it turns out that the quantity and price set by the monopolist at the equilibrium are the same over time. However, they are so for quite different reasons. In the first period, the monopolist increases quantities because it internalises the learning effect; in the second period, this externality is absent (it is the last period), but the equilibrium quantity increases because marginal costs have decreased.[10]

However, there is a sense in which learning makes the monopolist behave more competitively in the earlier periods. If one computes the market power (as proxied by the Lerner index: $L_t = (p_t^* - C'_t)/p_t^*$) exercised by the monopolist one obtains that market power is lower in the first period than in the second:

$$L_1 = \frac{(1-c)(1-l)}{1+c-l} < L_2 = \frac{(1-c)}{1+c-l}. \tag{24}$$

An introduction to elementary game theory

In this section I try to give a simple and short introduction to the most elementary concepts of game theory, that are also the ones that I use in most of the technical section of the books.[11]

Perhaps the simplest way to introduce the reader who is not familiar with game theory is to start with a simple example of a game, that I keep as abstract as possible so as not to distract attention with real world stories (I turn later to some more realistic applications).

Consider a game played between two *players*, call them player A and player B. Player A has to decide between two possible *actions*: a_1 or a_2; simultaneously, player B chooses among three possible actions b_1, b_2, or b_3. Each pair of actions will be associated with a certain *payoff* for each player, that is what a player receives after each of them has chosen an action. The pay-offs can be summarised in a payoff matrix such as in table 8.1. For instance, if player A chooses action a_1 and player B action b_2, that is, for the pair (a_1, b_2), the players' payoff is (2,5): A receives 2 and B receives 5. Assume also this is a game with *perfect information*, that is players have perfect knowledge of the actions available to each and of the pay-offs associated with them (this assumption is kept throughout unless otherwise indicated).

INSERT Table 8.1. A simple game

[10] These two effects exactly compensate each other in this model. See Tirole (1988: 72) for a learning model where output might even increase over time.

[11] Some sections marked with ** use concepts, such as Bayesian Nash equilibrium and sequential equilibrium, that are not explained here but are briefly introduced directly in those sections.

When one looks at a game, one would like to be able to predict its outcome. In this particular game, for instance, one would like to answer the question “if A and B were called to play this game, which actions will they choose?”, or, which is equivalent, “what will be the final result of the game?”.

It is clear that we could not reasonably expect that certain pairs of actions will be the result of the game. For instance, the pair (a_2, b_1) is unlikely to occur: if player B chose action b_1, player A would not respond with action a_2, that gives her the payoff 0, but rather with action a_1, that gives her the payoff 2. (Similarly, if player A chose a_2, player B would prefer to respond with action b_2, that gives him payoff 3 rather than the lower payoff 2 he would obtain by choosing action b_1 or action b_3.) In other words, we could not expect from two rational players that the final outcome of the game they play is a pair (a, b) such that a is not the best response to b and vice versa. For instance, in the game above, we would not expect (a_1, b_1) to be the final outcome of the game: if A plays a_1, B will prefer to play b_2, as this gives him a higher payoff (5, rather than 0).

2.3 Nash equilibrium

The concept of *Nash equilibrium* is based on such an idea: it predicts that the outcome of the game (the “equilibrium”) is given by the set of actions such that, for each player, each action is the best response to the actions of all other players. Equivalently, a set of actions represents a Nash equilibrium if none of the players has an incentive to deviate from its action given the actions of all other players.[12]

More formally, in a game with n players, denoting with A_i the set of actions available to player i (with $i = 1, ..., n$), and with $\pi_i(\alpha_1, \alpha_2, .., \alpha_i, .., \alpha_n)$ player i's payoff, the n-tuple $(\alpha_1^*, \alpha_2^*, .., \alpha_i^*, .., \alpha_n^*)$ is a Nash equilibrium if:

$$\pi_i(\alpha_1^*, \alpha_2^*, .., \alpha_i^*, .., \alpha_n^*) \geq \pi_i'(\alpha_1^*, \alpha_2^*, .., \alpha_i', .., \alpha_n^*), \text{ for all } i = 1, 2, .., n, \text{ and all } \alpha_i \in A_i. \tag{25}$$

In other words, $(\alpha_1^*, \alpha_2^*, .., \alpha_i^*, .., \alpha_n^*)$ is a Nash equilibrium if, for each player i, playing α_i^* is the best possible response given the other players play $(\alpha_1^*, \alpha_2^*, .., \alpha_{i-1}^*, \alpha_{i+1}^*, .., \alpha_n^*)$.

We can now return to the game described in figure 1, and look for the Nash equilibrium (or equilibria) of that game. To do so, let us write down the best responses for each player: the equilibrium, if it exists, will be given by the pair of actions that are mutually best responses.

Player A's best responses (indicated in bold) to each of B's actions are: $(\mathbf{a}_1, b_1), (\mathbf{a}_1, b_2), (\mathbf{a}_2, b_3)$.

[12] To keep things as simple as possible, the Nash equilibrium is defined here in terms of *actions*. It could also be defined in terms of strategies. A *strategy* s is a “rule” which tells a player which action to choose at any given time t of the game, for any given history of the game prior to t. Denoting with S_i the set of strategies available to player i (with $i = 1, ..., n$), and with $\pi_i(s_1, s_2, .., s_i, .., s_n)$ player i's payoff, the n-tuple $(s_1^*, s_2^*, .., s_i^*, .., s_n^*)$ is a Nash equilibrium if: $\pi_i(s_1^*, s_2^*, .., s_i^*, .., s_n^*) \geq \pi_i'(s_1^*, s_2^*, .., s_i', .., s_n^*)$, for all $i = 1, 2, .., n$, and all $s_i \in S_i$. In one-shot simultaneous games, the concepts of strategy and action coincide.

Player B's best responses (indicated in bold) to each of A's actions are: $(a_1, \mathbf{b}_2), (a_2, \mathbf{b}_2)$.

It is straightforward to see that (a_1, b_2) is the unique Nash equilibrium of this game, as it is the only pair of actions that are the best response to each other. Equivalently, it is the only pair of actions such that no player has an incentive to deviate from it, given the action of the rival.

2.3.1 The prisoners' dilemma

Perhaps the most famous game, since it has applications in several fields, from economics to politics, is the so-called prisoners' dilemma,[13] illustrated by the payoff matrix of table 8.2.

INSERT Table 8.2. The prisoners' dilemma game

This is a perfectly symmetric game that has only one Nash equilibrium, (low, low). (The reader can easily check by drawing the best responses for each player, as indicated above; or by verifying that no player would play *high* given that the other plays *low*, whereas any other pair of actions cannot be an equilibrium since a player would have an incentive to deviate from it.)

First, note that the game is very particular because each player has a *dominant strategy*: whether its rival plays *high* or *low*, a player always prefers to play *low*. Second, the outcome of the game is *Pareto inferior*: both players would be better off if they played $(high, high)$.

2.3.2 Coordination games and equilibrium selection

In the previous two games there was a unique Nash equilibrium. It is possible, though, that several equilibria co-exist, as in the game (a version of the so-called *battle of the sexes*) described in table 8.3.

INSERT Table 8.3. The battle of the sexes game

Anna and Bruno have decided to go to a restaurant together, but have not discussed which one, and cannot communicate with each other before dinner time. There are two good restaurants in town, one offers Indian and the other Thai food. Anna would prefer Indian and Bruno Thai, but for both the most important thing would be to end up in the same place. It is easy to check that this game has two Nash equilibria: $(Indian, Indian)$ and $(Thai, Thai)$.

It is far from uncommon to find games with multiple equilibria, and the problem in such cases is how to choose among them. In many economic applications, for instance, we are faced with such a situation, and unless we could refine our analysis the game would have scarce predictive power. This issue is addressed by *refinements* to the Nash equilibrium concept that try to select between different equilibria.

[13] So-called, because the original story behind this game is one where the players are two prisoners accused of a crime, and kept in separate cells, each of them having to decide whether to confess or deny the crime.

Pareto dominance as an equilibrium selection device Sometimes, as in the battle of the sexes game above, which is perfectly symmetric, even these refinements would not help much, but in others there might be one equilibrium that, for one reason or the other, might be more appealing and more likely for players to coordinate upon.[14]

Consider for instance a slightly different version of the battle of the sexes, as in table 8.4.

INSERT Table 8.4. A pure coordination game

Here there are again two equilibria: $(Indian, Indian)$, and $(Italian, Italian)$. However, ending up in the Italian restaurant would be "Pareto-inferior" for the players: that is, in plain words, both Anna and Bruno would have higher pay-offs if they had dinner in the Indian restaurant. Therefore, using Pareto-dominance as a criterion for selection among different equilibria, we would predict that Anna and Bruno will go to the Indian restaurant, a somehow reasonable prediction, and one which enjoys support from experimental evidence.[15,16]

Another criterion for equilibrium selection is the elimination of weakly-dominated strategies.

Elimination of weakly dominated strategies Consider two players A and B whose set of possible strategies is S_i (with $i = A, B$). We say that for player i a strategy $\underline{s_i}$ is weakly dominated if there exists a strategy s'_i such that $\pi_i(s'_i, s_j) \geq \pi_i(\underline{s_i}, s_j)$ for all s_j (with $i \neq j$), and there exists at least one strategy s_j for which this inequality holds with a strict sign.

Consider for instance the game illustrated in table 8.5:

INSERT Table 8.5. An asymmetric game

It is easy to check that this game has two equilibria: (p_L, p_L) and (p_H, p_H). The criterion of the elimination of weakly dominated strategies uniquely selects the equilibrium (p_H, p_H). Indeed, for player A we have $\pi_A(p_L, p_L) =$

[14] History or other circumstances might in some particular cases suggest that one equilibrium is more likely to be played than another. In the words of Schelling (1960), there might be *focal* points. For instance, if Anna and Bruno have so far always met at the Indian restaurant and never had dinner at the Thai restaurant, the former equilibrium is more likely than the latter. Some discussions on how certain prices might be focal can be found in chapter 4, where models having several, more or less collusive, equilibria, are analysed. There, I also discuss other elements that might resolve the uncertainty about the equilibrium, such as communication or first-mover advantages (Anna might move first, make a reservation at the Indian, and then make it known to Bruno). But of course all these are different games than the one-shot simultaneous moves game I am discussing here.

[15] Some experimental evidence is mentioned in chapter 4, that analyses repeated games that are very similar to the one described in this simple game, as they present a multitude of equilibria that can be Pareto-ranked.

[16] It is important to stress that here I am using Pareto-dominance to select between two Nash equilibria, whereas in the prisoners' dilemma game the pair giving the Pareto-superior outcome was not an equilibrium of the game, and therefore very unlikely to be the outcome of the game.

$\pi_A(p_L, p_H) = 0$ and $\pi_A(p_H, p_L) = 0 < \pi_A(p_H, p_H) = 2$, which implies that for A strategy p_L is weakly dominated by strategy p_H. Further, for player B we have $\pi_B(p_L, p_L) = \pi_B(p_L, p_H) = 0$ and $\pi_B(p_H, p_L) = -2 < \pi_B(p_H, p_H) = 0$, that is, for B strategy p_L is weakly dominated by strategy p_H.[17]

2.3.3 Mixed strategies

So far, I have restricted attention to so-called *pure strategies*: players decide whether to play a certain action or not but cannot randomise between them, that is they have to choose one action with probability one. Some games have no equilibrium in pure strategies, and it is easy to check that one such game is that described in table 8.6 (a version of the *matching pennies* game).

INSERT Table 8.6. The matching pennies game

Mixed strategies However, one might want to consider the possibility that players use *mixed strategies*, that is, that they have a probability distribution over their pure strategies (in other words, players randomise among pure strategies).[18] For instance, in the matching pennies game of table 8.6, player A might play *heads* with a certain probability p, *tails* with probability $1-p$; and player B might play *heads* with a probability q, *tails* with probability $1-q$. The concept of Nash equilibrium and of best responses must then be re-cast in terms of mixed strategies. Note that for a mixed strategy to make sense, it must be that a player is indifferent between two (or possibly more) pure strategies, that is, they should give her the same payoff given the (mixed) strategy of the other player(s).

In the game of table 8.6, for instance, the equilibrium in mixed strategies is the one where $p = q = 1/2$. To see why, look first at the payoff player A makes given the mixed strategy of player B. If A chooses *heads*, she gets: $\pi_A(heads) = -1(q)+1(1-q) = 1-2q$. If she chooses *tails*, she gets: $\pi_A(tails) = 1(q)-1(1-q) = -1+2q$. To be willing to randomise between *heads* and *tails*, player A must be indifferent between the two, that is it must be $\pi_A(tails) = \pi_A(heads)$. It is easy to check that this is true for $q = 1/2$.

A similar easy calculation (here everything is symmetric) shows that player B is indifferent between *heads* and *tails* (that is, he is willing to randomise between the two) if player A chooses *heads* with probability $p = 1/2$ and *tails* with probability $1-p = 1/2$.

[17] Note that in this example Pareto dominance would also select the same (p_H, p_H) equilibrium pair.

[18] If it sounds strange that a player "throws dice" to decide which action to take, a more appealing interpretation of mixed strategies exists: "the crucial feature of a mixed-strategy Nash equilibrium is not that player j chooses a strategy randomly, but rather that player i is uncertain about player j's choice; this uncertainty can arise either because of randomization or (more plausibly) because of a little incomplete information." Gibbons (1997: 140). Gibbons (1997: 138-140) shows how to interpret mixed strategies using games with *incomplete information*, that is where at least one player is not sure about the other player's payoff.

Nash (1950) has showed that any game with a finite number of players and where each player has a finite number of pure strategies, has a Nash equilibrium (in mixed strategies if no equilibrium in pure strategies exists).

Note also that equilibria in mixed strategies can co-exist in the same game with equilibria in pure strategies. For instance, the reader can check that the battle of the sexes game of table 8.3 admits one Nash equilibrium in mixed strategies where both A and B chooses ($Indian$) with probability 1/3 and A chooses ($Thai$) with probability 1/2.

2.4 Dynamic games, and sub-game perfect Nash equilibrium

So far we have just looked at static games, where players choose their actions simultaneously. Dynamic games are games where players move sequentially or they move more than once. Such a game might be as follows. First, player 1 chooses an action a_1 from the set of her feasible actions; second, player 2 chooses an action a_2 from the set of his feasible actions; after both have played, players 1 and 2 receive the pay-offs associated with the pair (a_1, a_2).

An example will help illustrate why the concept of Nash equilibrium needs to be refined to deal with dynamic games. Consider two players: firm I is an incumbent firm, firm E is a potential entrant in the industry. First, firm E decides whether to enter or not in the market; second, firm I decides whether to *accommodate* entry (for instance, setting a high price) or to $fight$ entry (setting a low price). Table 8.7 represents this game in the so-called *normal form*, that is with the usual payoff matrix.

INSERT Table 8.7. The entry deterrence game

This game has two Nash equilibria: (*enter*, *accommodate*) and (*stay out*, $fight$). However, the second equilibrium, (*stay out*, $fight$), is an unlikely outcome of this game. At this equilibrium, firm E chooses to stay out because firm I would fight entry, thereby making entry unprofitable to E. However, the threat of fighting entry by firm I if entry occurs is not credible. Indeed, if entry did take place, firm I would rather accommodate it, since it would get a higher payoff than if it fought it.

We therefore need a refinement of the Nash equilibrium concept, that allows to rule out equilibria - such as the one described above - that are based on non-credible threats. In other words, we need to ensure that each player always plays optimally in each situation, even in those that are not along the equilibrium path. (In the game above, an equilibrium is given by (*stay out*, $fight$), but $fight$ is not optimal if the game reached a point where $entry$ occurs.)

The concept of *sub-game perfect Nash equilibrium* (SPNE) corresponds to this requirement. It is the set of strategies for each player such that the strategies form a Nash equilibrium in any sub-game of the game (and not only along the equilibrium path), that is any sub-set of the game that starts from any point at which the complete history of the game prior to that point is common knowledge for all the players (each player knows it, and knows that the other knows it...).

To see how to look for the SPNE of a game, it is useful to describe a game in its *extensive form* (or game tree). For the game above, this is done in Figure 8.1.

INSERT Figure 8.1. Extensive form of the entry deterrence game

The game tree of figure 8.1 contains two sub-games: the whole game; and the sub-set of the game which starts when player I is about to move after player E has played *enter*. What the SPNE equilibrium concept requires is then that all players play optimally at each sub-game.

2.4.1 Backward induction

To find the (pure strategy) SPNE equilibrium, we have to start from the last branches of the tree and move backwards. At the sub-game starting after player E has chosen *enter*, player I will choose *accommodate* (it gets 5 rather than 0). We can now move backwards to the first sub-game, which is the initial node of the game. Player E knows that if it chooses to stay out it will get 0 whatever I does, and anticipates correctly that if it enters player I will accommodate and it will make a profit of 4. Therefore, it will enter: $(enter, accommodate)$ is the only SPNE of this game.

Note that SPNE is used also in games of "almost perfect information", where more than one player moves in the same sub-game. Consider for instance the game depicted in Figure 8.2. This is a game where player E decides first whether to *stay out* or *enter*; after observing E's decision, active firms decide (simultaneously, if E has entered) whether to sell *low* or *high* quality.

INSERT Figure 8.2. Extensive form of a quality game

Figure 8.2 represents the game in its extensive form. The "oval" around E's nodes is called an *information set*: it is a convention to represent the idea that when E decides it does not know at which node it is, that is whether I has chosen *low* or *high* quality (this can be either because I and E take decisions on quality simultaneously, as I have said above, or because I takes its choice before but E does not observe I's choice before its own choice: the two are equivalent).

Note that this game has three sub-games: the whole game; the sub-set of the game which starts when player I is about to move after player E has played *stay out*; the sub-set of the game which starts when player I is about to move after player E has played *enter*. (The nodes starting after I has played are not sub-games because the complete history of the game prior to that point is not common knowledge for all the players: at the moment of choosing quality, player E does not know what I has played.)

By using backward induction, it is easy to find that there are two pure strategy SPNE of this game: $(enter, low, high)$ and $(enter, high, low)$.[19]

[19] If E enters, the last stage of the game has two equilibria: $(low, high)$ and $(high, low)$. Find first the equilibrium of the whole game if $(low, high)$ is played;then,if $(high, low)$ is played.

Another type of game that belong to the category of games of "almost perfect information", and that can be solved by backward induction, are those where the same game is repeated finitely many times. For instance, the prisoners' dilemma game repeated over and over, but for a finite number of times is one such game (see also section 3.2.2). However, backward induction cannot be used for *infinite horizon* games: backward induction requires starting from the final node of the game and moving backwards, whereas no such final node can be found if the game lasts forever.

The concept of SPNE, and the method of backward induction to look for it, can be found everywhere in this book since it is used to solve dynamic games that are widely used in industrial organisation theory (see also below in this section). Examples of a game being played for infinite number of times is at the core of chapter 4 on collusion (see also section 3.2.3 below).

3 Oligopoly I: Market competition in static games

This section introduces the reader to simple models of product market competition in static games, and under exogenously given investment choices of firms (section 4 will consider games where firms decide their investments in capacities, R&D, advertising, quality or product positioning). Section 3.1 also assumes that the firms produce a perfectly homogenous good, and examines the different benchmark cases of product market competition: price competition (the *Bertrand* model), quantity competition (the *Cournot* model) and joint-profit maximisation. Section 3.2 assumes that the firms sell differentiated products, and again derives the equilibrium outcomes for the three benchmark cases of price competition, quantity competition and joint-profit maximisation. Throughout these two sections, we shall use the Nash equilibrium concept introduced in the previous section: for given characteristics of the goods (that is, whether they are homogenous or differentiated), the firms play a one-shot game where their actions might consist of prices or quantities, and the Nash equilibrium is the appropriate concept to study the outcome of the game.

Section 3.2.1 will consider a repeated game of product market interaction: again, firms have exogenously given product characteristics, but they meet period after period in the marketplace. For those games, the relevant equilibrium concept is the one of sub-game perfect Nash equilibrium.

3.1 Product market competition with homogenous goods

In this section I assume that the firms' products are homogenous, that is are seen as perfect substitutes by consumers. I will analyse in turn the three benchmark cases where firms compete in prices, in quantities, and where they choose their actions so as to maximise joint profits.

3.1.1 Price competition (Bertrand model)

Consider two firms (but the extension to n firms would be straightforward and offer the same results) that:

A1) sell *homogenous* goods;

A2) play a *one-shot game*;

A3) independently and simultaneously choose the *price* at which they want to sell their product;

A4) have *no capacity constraints*, that is, they are able to serve all demand that is addressed to them;

A5) have the same *identical marginal cost*, c, and no fixed costs.

Consumers address the firms according to the following demand function:

$$D_i(p_i, p_j) = \begin{cases} D(p_i), & \text{if } p_i < p_j \\ D(p_i)/2, & \text{if } p_i = p_j \\ 0, & \text{if } p_i > p_j \end{cases}, \tag{26}$$

that is: if a firm sets a price lower than its rival, then all consumers will address it (and vice versa: nobody addresses the firm setting the higher price); if both firms set the same price, consumers are indifferent between one or the other, and it is assumed that they equally split their demand between the two firms.

This being a one-shot game, the appropriate solution concept is the Nash equilibrium. In this model, a Nash equilibrium in prices, or Bertrand equilibrium, is a pair of prices (p_i^*, p_j^*) such that $\pi_i(p_i^*, p_j^*) \geq \pi_i'(p_i, p_j^*)$, for any $i, j = 1, 2$, with $i \neq j$ and any p_i in the real numbers. I am going to show the following result:

(Bertrand equilibrium) The unique price equilibrium of this game is given by $p_i^* = p_j^* = c$, with $\pi_i(p_i^*, p_j^*) = \pi_j(p_i^*, p_j^*) = 0$.

Proof. To prove this result, we need to prove first that the proposed one is a Nash equilibrium of the game, and then that its is the only one.

Step 1. To see that $p_1^* = p_2^* = c$ is a Nash equilibrium is straightforward. To be a Nash equilibrium, no firm must have incentives to deviate from it given that the other plays according to the equilibrium strategy. Suppose then that $p_1^* = c$, does firm 2 have an incentive to deviate and sell at a different price than $p_2^* = c$? By selling at marginal cost, firm 2 serves half of the market but makes zero profit. If it deviates and sets a *lower* price than c, it will serve all the market, but make losses; if it deviates and sets a *higher* price, no consumers will address it, and therefore will make zero profits, thus not improving its situation relative to playing the candidate equilibrium. Therefore, firm 2 will have no incentive to deviate. Likewise, given the perfect symmetry between the two firms, firm 1 has no incentives to deviate.

Step 2. Let us reason by contradiction. Suppose that there is a different candidate equilibrium, and show there is at least one deviation that would make better off a firm, thereby breaking the candidate equilibrium.

- $p_i^* = p_j^* = p^* > c$: at this candidate equilibrium, firms share the market equally and make positive profits $\pi(p^*, p^*) = (p^* - c)D(p^*)/2$. However, given the price p^* of firm i, its rival firm j can set a price p_j' which is a shade less than p^* (that is, $p_j' = p^* - \varepsilon$) and get the whole market. It would therefore make a profit $\pi_j'(p_i^*, p_j') = (p^* - c - \varepsilon)D(p^* - \varepsilon)$ that, for ε small enough, is clearly higher than $\pi(p^*, p^*)$ (since a marginal reduction in the price results in a disproportionate increase in demand, which doubles). Therefore, this cannot be a Nash equilibrium of the game.

- $p_i^* = p_j^* = p^* < c$: at the candidate equilibrium, both firms make losses. Trivially, this cannot be an equilibrium since a firm would have an incentive to deviate and set a price at or above marginal cost: it would not sell anything and therefore make zero profits (better than negative profits).

- $p_i^* > p_j^* \geq c$: if $p_j^* = c$, firm j makes zero profits, but given that firm i sets a price above cost, it can improve its position by charging any price between c and firm i's price, since it would still get the whole demand but would command a positive margin. In particular, the optimal deviation would be to set the price $p_i^* - \varepsilon$ and get profits $(p_i^* - c - \varepsilon)D(p_i^* - \varepsilon) > 0$. Therefore, this candidate equilibrium, and by similar reasoning all pairs with asymmetric prices, cannot be a Nash equilibrium of the game. If $p_j^* > c$, firm i would also have an incentive to deviate, and charge a price just below p_j^*.

■

The Bertrand result is a striking one. Despite the fact that there are only two firms in the industry, they will end up selling at marginal cost, and get zero profits. As we shall see below, this is not a robust result, as it crucially depends on a series of strong assumptions: by relaxing, in turn, each of assumptions A1)-A5), one obtains equilibria where prices are above marginal cost and firms make positive profits. Nevertheless, this case provides a useful benchmark, corresponding to the lower bound that equilibrium prices can take. In other words, the Bertrand equilibrium corresponds to the toughest possible degree of product market competition.

Before introducing the other main benchmark case, that of Cournot competition, let me briefly describe the Bertrand game in two interesting cases: (1) under asymmetric firms; (2) under capacity constraints.

Bertrand under cost asymmetric firms Consider exactly the same game as above, but relax assumption A5) and assume that the two firms have asymmetric costs. Firm 1 and 2 have respectively marginal cost c_1 and c_2 with $c_1 < c_2$. There are two possible solutions to this game. In the first case, firm 1 is so much more efficient than firm 2 (firm 1's costs are much lower), that it can behave as if it was a monopolist: firm 2's cost is so high that firm 1's monopoly price, p_1^m is below c_2. In the second case, if the firms' costs are close enough to each other (how close is to be specified below), then firm 1 will set a price slightly below firm 2's marginal cost and get the whole market.

To keep the analysis simple, I develop these two cases under the assumption that firm i's demand function is given by $D_i = 1 - p_i$ if $p_i < p_j$, $D_i = 0$ if $p_i > p_j$, and $D_i = (1 - p_i)/2$ if $p_i = p_j$.

Large asymmetries Consider the case where firm 2's costs are large enough relative to firm 1's (more precisely, $c_2 \geq (1 + c_1)/2$, as will be seen below). If firm 1 were alone in the market, it would choose its price so as to maximise $\pi = (1 - p)(p - c_1)$. We have seen that the monopolist's problem is easily solved by taking $d\pi/dp = 0$ and re-arranging, which gives as result:

$$p_1^m = \frac{1 + c_1}{2}. \tag{27}$$

As long as p_1^m is below the marginal cost of firm 2, firm 1 can charge the monopoly price and get the whole market undisturbed, as firm 2 would have no incentive to undercut it: if it did set a price below p_1^m, all consumers would address it and it would accumulate losses. Better then to charge a price higher than p_1^m and get zero profits.

Therefore, if $c_2 \geq (1 + c_1)/2$, at equilibrium firm 1 will set $p_1^m = (1 + c_1)/2$, firm 2 a price $p > p_1^m$. Firm 1 will get monopoly profits $(1 - c_1)^2/4$, and firm 2 zero profit.

In the study of innovations, this case corresponds to the case where a firm gets a *drastic innovation*, that is an innovation that makes it so much more competitive than its rival, that it can behave as a monopolist (see chapter 2 for some applications of this case).

Small asymmetries Consider now the case where the firms' costs are close enough: $c_1 < c_2 < (1 + c_1)/2$. In this case, firm 1 setting its monopoly price cannot be the equilibrium of the game: if firm 1 sets $p_1^m = (1+c_1)/2$, firm 2 will charge a price $p = (1 + c_1)/2 - \varepsilon$, get all the market, and make a positive profit (on each unit sold, it will make a gain equal to $(1 + c_1)/2 - \varepsilon - c_2$).

The following is instead a Nash equilibrium of the game: $(p_1^*, p_2^*) = (c_2 - \varepsilon, c_2)$, that is, firm 2 charges a price equal to marginal cost and firm 1 a price which is a shade below that. It is easy to check that firm 2 has no incentive to deviate from this pair, as undercutting $p_1^* = c_2 - \varepsilon$ would leave it with losses rather than zero profits, whereas increasing the price above c_2 leaves it again with zero profits. Firm 1 has no incentive to deviate either: by setting a higher price, say $p_1 = c_2$ it would have to share the market with the rival (or lose it completely to firm 2 if the price was above c_2),[20] whereas by setting a lower price it would continue to get all the market but it would sell at a lower margin, and thus making lower profits.

[20] There is a very minor technical detail here, which is that firm 1 wants to choose ε as close as possible to zero, but of course for any given small number ε it is always possible to find another number smaller than it. To resolve this technical problem, an *escamotage* is sometimes used, by assuming that for identical prices all the demand goes to the firm that has lower costs. The equilibrium of the game then becomes $(p_1^*, p_2^*) = (c_2, c_2)$, with all demand going to firm 1.

At this equilibrium, firm 1 makes profits $(c_1 - c_2)(1 - c_2)$, and firm 2 gains zero.

In principle, there exist other equilibria of this game, but they are less "reasonable". Consider for instance a price $p \in (c_1, c_2)$. It is easy to see that the pair $(p_1^{**}, p_2^{**}) = (p - \varepsilon, p)$ represents an equilibrium of the game.

However, such an equilibrium looks less appealing than the equilibrium (p_1^*, p_2^*) identified above, and indeed the main criteria of equilibrium selection would select the latter over the former. Under Pareto-dominance, for instance, (p_1^*, p_2^*) would be chosen because it gives higher profits for firm 1 while keeping the same profits (zero) for firm 2. Elimination of weakly dominated strategies also selects (p_1^*, p_2^*) as the only equilibrium of the game. To see this, note that when player 2 chooses action $p_2^* = c_2$ it gets $\pi_2(c_2 - \varepsilon, c_2) = 0$ and $\pi_2(p - \varepsilon, c_2) = 0$, where $p \in (c_1, c_2)$; when it chooses action $p_2^{**} = p$, it gets $\pi_2(p - \varepsilon, p) = 0$ if player 1 chooses $p - \varepsilon$, but it gets $\pi_2(p + \varepsilon, p) < 0$ if player 1 chooses $p + \varepsilon$. The equilibrium (p_1^{**}, p_2^{**}) contains a strategy that is weakly dominated for player 2.

Bertrand under capacity constraints Let us now go back to the symmetric case treated initially, but relax now assumption A4), and assume instead that each firm holds a capacity $k_i < D(p_i = c)$: when charging at marginal cost, a firm would have to supply a larger number of units than its capacity would allow it to do. It is then easy to see that under this new assumption $p_i^* = p_j^* = c$ is not a Nash equilibrium of the game any longer:

Remark 1 *If $k_i < D(p_i = c)$, $(p_i^*, p_j^*) = (c, c)$ is not a Nash equilibrium of the game.*

Proof. To prove this result, we just need to find a deviation that leaves one of the firms better off. Now, at the candidate equilibrium $p_i^* = p_j^* = c$, firm i makes zero profits. But if it deviates and sells at a price $p_1' > p_j^* = c$, some consumers will address it (at least as long as p_1' is not too high) and it will make positive profits. This is because all consumers would like to buy from firm j, that sells at a lower price, but j cannot serve them all (as its capacity $k_j < D(c)$). Some consumers are rationed, and will have to buy instead from firm 1, that will therefore make positive profits (it will sell a positive number of units at a positive margin). ■

I limit myself here to this proof that under capacity constraints the Bertrand result does not hold. To find the equilibrium of the price competition game under capacity constraints requires specifying a *rationing rule*, that is a rule that allocates consumers between the firms (some cannot buy from the low price firm). It also turns out that an equilibrium in pure strategies does not exist if firms have large enough capacities. Since throughout the book I never deal with capacity constrained models, it would be an unnecessary complication to carry out the complete analysis of the capacity-constrained price competition game. I refer the interested reader to Kreps and Scheinkman (19xx): in a two-stage game where firms (simultaneously) first choose capacities and then prices, the final equilibrium outcome is the same as in a one-shot game where firms choose quantities. In other words, the Cournot equilibrium could be interpreted not only literally, that is as the outcome of a game where firms choose the output they bring to the market, but also as the outcome of a game where firms choose their capacities and then their prices.[21]

3.1.2 Quantity competition (Cournot model)

Consider now the same one-shot game analysed in section 3.1.1, but relax assumption A3) and assume instead that firms choose *quantities* they want to bring to the market, rather than prices. For simplicity, let us look first at the symmetric case, where both firms have the same marginal cost $c < 1$, and they face a linear demand function $p = 1 - Q$, with $Q = q_1 + q_2$ being total industry output.[22]

Firm i's profit is given by $\pi_i = p(q_i, q_j)q_i - cq_i$. The first step to identify the Nash equilibrium of this game is to look for the best reply function of a firm for any given quantity chosen by its rival. Taking q_j as given, firm i solves therefore the following problem (for $i, j = 1, 2$, and $i \neq j$):

$$\max_{q_i} \pi_i(q_i, q_j) = (1 - q_i - q_j - c)q_i, \text{ given } q_j. \tag{28}$$

This problem is solved by taking $d\pi_i(q_i, q_j)/dq_i = 0$, that is:

$$1 - 2q_i - q_j - c = 0. \tag{29}$$

The previous expressions ($1-2q_1-q_2-c = 0$ for firm 1, and $1-2q_2-q_1-c = 0$ for firm 2) implicitly represent the *reaction functions* (or *best reply,* or *best response functions*) of each firm, whose meaning we have already discussed in section 2.2.3. It is convenient to represent the reaction functions in the same plane (q_1, q_2). To this end, let us write them as:

$$R_1 : \ q_2 = 1 - 2q_1 - c; \ R_2 : \ q_2 = \frac{1 - q_1 - c}{2}. \tag{30}$$

Figure 8.3 illustrates the two reaction functions.

INSERT Figure 8.3. Reaction functions in the Cournot model

Note that they are negatively sloped: the higher firm j's output q_j the lower firm i's best reply output. The slope of the reaction functions carries important effects when analysing dynamic models, as we shall see in section 4.1 below.

[21] See also Maggi (1996), where constraints are 'soft', as a firm can increase its output beyond capacity by incurring an additional variable cost. Maggi's formulation is simpler as it allows us to find an equilibrium in pure strategies of the price game independently of the capacity levels. Note, however, that the Kreps-Scheinkman result applies only when using the efficient-rationing allocation rule, as showed by Davidson and Deneckere (19xx).

[22] The assumption $c < 1$ serves to guarantee the viability of the market. Otherwise, there is no price at which firms would be willing to supply demand as they would get negative profits.

The figure also illustrates the isoprofit functions of each firm, that is the locus of the points such that different values of q_1, q_2 give the same value of profit to a firm, $\pi_i(q_i, q_j) = k$. Note that lower isoprofit curves for firm 1, and isoprofit curves more to the left for firm 2, are associated with higher profits (given the same q_i, a lower q_j would increase prices and profits of firm i).[23]

As we know from the discussion of the game theory section, the Nash equilibrium of the game will be determined by the intersection of the two reaction functions, that elementary algebra shows to be:

$$q_1^C = q_2^C = \frac{1-c}{3}, \tag{31}$$

the label "C" standing for Cournot.[24] After substituting, one obtains also prices and profits as:

$$p^C = \frac{1+2c}{3}, \quad \pi_1^C = \pi_2^C = \frac{(1-c)^2}{9}. \tag{32}$$

Note that at the Cournot equilibrium the market price is above marginal costs (recall that $c < 1$, so $(1 + 2c)/3 > c$) and firms make positive profits. The fact that at an equilibrium under quantity competition firms set a higher price than under a price competition equilibrium descends from the very concept of Nash equilibrium. In a Bertrand game, a firm has to choose its optimal price by taking as given the price of the rival. By undercutting the latter, it will get all the demand, which gives a very strong incentive to decrease one's price. In a Cournot game, instead, a firm chooses its quantity given the quantity of the rival. Therefore, an expansion of one's output might allow one to capture a higher market share but does not allow one to capture all demand precisely because the rival's output is taken as given. The incentives to compete aggressively are accordingly weaker than in the price competition game. In other words, price competition is tougher than quantity competition and other things being equal, firms' prices and profits are lower. This result will be confirmed also when looking at the differentiated products case.

[23] Note also that an isoprofit function of firm 1 must be tangent to the horizontal line when it crosses firm 1's best reply function, and an isoprofit function of firm 2 must be tangent to the vertical line when it crosses 2's best reply function. To see this, write a generic isoprofit function of firm 1 as $\pi_1(q_1, q_2) = (1 - q_1 - q_2 - c)q_1 = k$. To find its slope, write the total differential as $d\pi_1 = (1 - 2q_1 - q_2 - c)dq_1 - q_1 dq_2 = 0$. Therefore, its slope must be $dq_2/dq_1 = (1-2q_1-q_2-c)/q_1$, whose numerator is nil when the FOC is satisfied (that is, along the reaction function). Similarly, write a generic isoprofit function of firm 2 as $\pi_2(q_1, q_2) = (1 - q_1 - q_2 - c)q_2 = z$. The total differential is $d\pi_2 = (1 - q_1 - 2q_2 - c)dq_2 - q_1 dq_1 = 0$. Its slope must be $dq_2/dq_1 = q_1/(1 - q_1 - 2q_2 - c)$, whose denominator is zero when the FOC is satisfied: therefore, along the reaction function the slope is infinity.

[24] I focus here on the symmetric equilibrium of the game. There also exist two possible asymmetric Nash equilibria, where a firm, say 1, produces zero output and the other firm produces such a large output that if firm 1 brought even a very small quantity to the market the price would fall below marginal costs (thus giving it no incentive to deviate from zero output). This case is considered in detail in chapter 7.3.3.1 (on interoperability) and omitted here to save space.

Cournot with cost asymmetries Consider the same Cournot game as before, but assume that firms differ in their marginal costs, with $c_1 < c_2$. Each firm i $\max_{q_i} \pi_i(q_i, q_j) = (1 - q_i - q_j - c_i)q_i$, given q_j. The reader can check that the reaction functions are given by $q_i = (1 - q_j - c_i)/2$, and that the Cournot equilibrium quantities and profits (determined by the intersection of the reaction functions and substitution) are given by:

$$q_i^* = \frac{1 - 2c_i + c_j}{3}; \quad \pi_i^* = \frac{(1 - 2c_i + c_j)^2}{9}. \tag{33}$$

However, note that this solution holds only if costs are close enough to each other. Indeed, the equilibrium output is meaningful only if $q_2^* \geq 0$, or $c_2 \leq (1 + c_1)/2$. Otherwise, similarly to the case of Bertrand competition, firm 2 is so much less efficient than firm 1 that the latter can set, undisturbedly, its monopoly output and make monopoly profits.

Cournot with n firms Consider now again the case of symmetric cost firms, but extend the base model to n firms. For the i-th firm, the problem is:

$$\max_{q_i} \pi_i(q_1, .., q_i, .., q_n) = (1 - q_i - \sum_{j \neq i} q_j - c)q_i, \text{ given } q_j. \tag{34}$$

The FOC is given by $d\pi_i(q_1, .., q_i, .., q_n)/dq_i = 0$, that is:

$$1 - 2q_i - \sum_{j \neq i} q_j - c = 0. \tag{35}$$

At the symmetric equilibrium, $q_i = q_j$, and the FOC simplifies to $1 - 2q_i - (n - 1)q_i - c = 0$. It is then immediate to derive the equilibrium outputs, and by substitution the equilibrium price and profits, as:

$$q^* = \frac{1 - c}{1 + n}; \quad p^* = c + \frac{1 - c}{1 + n}; \quad \pi^* = \left(\frac{1 - c}{1 + n}\right)^2. \tag{36}$$

One can check that for $n = 2$ the equilibrium values in (36) correspond to the ones found above in expressions (31) and (32). But these results are interesting because they allow to study how the Cournot equilibrium outcome changes with n. In particular, it is easy to see that the larger the number of firms in the industry the closer one gets to the Bertrand outcome: $\lim_{n \to \infty} p^* = c$.

The Cournot model allows us therefore to capture the intuitive result that the more firms that co-exist in the industry the stronger competition will be in it, ranging from the monopoly outcome corresponding to $n = 1$ to the Bertrand outcome ($p = c$) when n tends to infinity. Such a result did not arise in the Bertrand model, where already with $n = 2$ one obtains that firms charge at marginal cost, a result which holds independently of the number $n \geq 2$ of firms in the industry.

3.1.3 Joint-profit maximisation

The final benchmark case of product market competition corresponds to the case where the oligopolists maximise joint profits, that is they behave as if they were a single firm. This case therefore corresponds to that analysed in section 2.2. As an illustration, consider n perfectly symmetric firms. Joint-profit maximisation implies:

$$\max_{q_1,..,q_i,..,q_n} \Pi = \sum_{i=1}^{n} \pi_i(q_1,..,q_i,..,q_n) = \sum_{i=1}^{n}(1 - q_i - \sum_{j \neq i} q_j - c)q_i. \tag{37}$$

In this particular case where firms are perfectly symmetric, this problem amounts to $\max_Q(1 - Q - c)Q$. The FOC is given by $d\Pi/dQ = 0$, that is $1 - 2Q - c = 0$, which results in the following equilibrium levels:

$$Q^M = \frac{1-c}{2};\ p^M = \frac{1+c}{2};\ \Pi^M = \frac{(1-c)^2}{4}. \tag{38}$$

Obviously, per-firm outputs and profits can be obtained by dividing by n:

$$q^M = \frac{1-c}{2n};\ \pi^M = \frac{(1-c)^2}{4n}. \tag{39}$$

The symmetric case is extremely simple because it is natural to assume that outputs and profits are shared equally, but the treatment of joint-profit maximisation becomes more complex when asymmetries among the firms exist.[25]

Note also that the joint-profit maximisation case should be seen as a benchmark case, that is as the limit of a situation where competition among the firms in the product market is very weak. For a treatment of collusion, that is how and whether firms are able to sustain a joint-profit maximisation outcome, see section 3.2.1 and, above all, chapter 4.

3.1.4 Benchmark models of product competition: A comparison

As a summary of the different cases analysed here as benchmarks of product market competition, it might be helpful to illustrates the equilibrium price as a function of the number n of firms in the industry, under the hypotheses that firms compete in prices, in quantities and that they behave as if they were a monopolist. Figure 8.4 does this, and shows for any given number of firms how the toughness of product market competition - or *toughness of price competition*, as Sutton (1991) calls it - varies across these cases, being highest under Bertrand (equilibrium prices are the lowest) and lowest under joint-profit maximisation (equilibrium prices are the highest).

INSERT Figure 8.4. Toughness of product market competition: Bertrand (B), Cournot (C), and joint-profit maximisation (M)

[25] I do not deal with this issue because it does not arise anywhere in the book.

The three cases analysed are useful benchmarks when discussing oligopoly issues. It would be difficult to say when in the real world one should expect product market competition to take one form or another. Joint-profit maximisation might perhaps correspond to a (nowadays very rare) situation where antitrust enforcement is so weak that a cartel can be easily sustained. Cournot competition might be associated with industries where firms cannot easily adjust their capacities (that is, when they choose prices after they have committed to a certain capacity or output), and Bertrand competition with industries where firms are instead not much constrained by capacities; that is, where they can easily adjust output in response to the demand they receive.[26]

3.2 Product market competition with (exogenously) differentiated goods

This section studies (one-shot) product market competition within two (non-spatial) differentiated good models. This amounts to relaxing assumption A2) of section 3.1.1, and again this will result in an equilibrium price above marginal cost (except in limiting cases). Throughout this section firms' products choices are taken as given.

3.2.1 A linear demand model of differentiated goods

Following Singh and Vives (1983), consider two firms, 1 and 2, that sell two products. Denote q_1 and q_2 the quantities of each good. There exists a continuum of consumers of the same type in the economy, each having the following utility function:

$$V = y + U(q_1, q_2), \tag{40}$$

where the linearity in the composite good y avoids income effects and provides a rationale for a partial equilibrium analysis of the differentiated sector. Indeed, the consumer's problem is $\max_{q_1,q_2,y} V$ subject to the budget constraint $p_1q_1 + p_2q_2 + p_yy = R$. To solve this problem, write the Lagrangean:

$$L = y + U(q_1, q_2) + \lambda(R - p_1q_1 - p_2q_2 - p_yy). \tag{41}$$

The FOCs are given by:

$$\begin{cases} \frac{\partial L}{\partial q_i} = \frac{\partial U(q_i,q_j)}{\partial q_i} - \lambda p_i = 0, \quad i = 1, 2; i \neq j, \\ \frac{\partial L}{\partial y} = 1 - \lambda p_y = 0, \\ \frac{\partial L}{\partial \lambda} = R - p_1q_1 - p_2q_2 - p_y = 0. \end{cases} \tag{42}$$

[26] Most manufacturing industries will probably be closer to Cournot than Bertrand. Industries where procurement is important might be an example of Bertrand competition: first a firm selects price and gets a contract and then it has to fulfill the order within a certain time. Sectors where it is costless and immediate to ship a good would also resemble Bertrand competition (think for instance of the music industry, where once the original track is recorded, making an additional CD is easy and cheap).

By taking the composite good as the numéraire, $p_y = 1$, one obtains $\lambda = 1$. Therefore, the FOCs relative to the differentiated good market become $\partial U(q_i, q_j)/\partial q_i - p_i = 0$, and can be analysed independently of the market for the composite good.

Once established that the quasi-linearity in the utility function V justifies a partial equilibrium analysis of the differentiated good market, let us use a specific functional form for the utility function. In particular, assume:

$$U(q_1, q_2) = \alpha q_1 + \alpha q_2 - \frac{1}{2}\left(\beta q_1^2 + \beta q_2^2 + 2\gamma q_1 q_2\right), \tag{43}$$

with the parameters used in the utility function satisfying the following assumptions (for $i = 1, 2,\ i \neq j$): (i) $\alpha > 0$, (ii) $\beta > 0$, (iii) $\beta > \gamma$. Assumption (iii) guarantees that the demand functions can be inverted, are of the right sign, and that there is a positive intercept in direct demands (see below).

Parameter γ indicates whether goods 1 and 2 are substitutes, independent or complements (and to what degree). As one can see from the utility function, if $\gamma > 0$ consuming the two goods together diminishes the consumer's utility (that is, they are substitutes); if $\gamma < 0$ consuming the two goods together increases the consumer's utility (that is, they are complements); if $\gamma = 0$ consuming the two goods together does not affect the consumer's utility (that is, they are independent).

The FOCs of the consumer problem are given by $\partial U(q_i, q_j)/\partial q_i - p_i = 0$, which become:

$$p_i = \alpha - \beta q_i - \gamma q_j, \quad i = 1, 2, i \neq j. \tag{44}$$

This is the system of inverse demands that can be used to study quantity competition. An inspection of these demand functions immediately suggests further interpretation of the parameters. The closer γ to β the more substitutes the two goods will be, with the case of perfect substitutes arising as the limit case $\gamma \to \beta$. One could therefore define an inverse measure of product differentiation as γ/β. This index takes values in $[0, 1)$, and attains its minimum value when the goods are independent (for $\gamma = 0$, that is when they are maximally differentiated), and its highest value when they tend to be perfect substitutes (for $\gamma \to \beta$).

By inverting the two expressions in (44) one obtains the system of direct demand functions, as:

$$q_i = a - bp_i + gp_j, \quad i = 1, 2, i \neq j, \tag{45}$$

where a, b, g satisfy:

$$a = \frac{\alpha(\beta - \gamma)}{\beta^2 - \gamma^2};\ b = \frac{\beta}{\beta^2 - \gamma^2};\ g = \frac{\gamma}{\beta^2 - \gamma^2}. \tag{46}$$

Note, therefore, that this is the linear demand function already used in section 2.2. Armed with the inverse and direct demand functions, we can now derive the equilibrium solutions under the usual benchmark cases of product market competition. For simplicity, assume that the two firms have the same constant marginal cost c and assume that $c = 0$ without loss of generality.

Quantity competition Consider first the case where firms compete in quantities. Each firm $i = 1, 2$ chooses q_i so as to maximise its profits $\pi_i = p_i(q_i, q_j)q_i$ for any given q_j, where $p_i(q_i, q_j) = \alpha - \beta q_i - \gamma q_j$. The FOCs are:

$$\frac{d\pi_i}{dq_i} = \alpha - 2\beta q_i - \gamma q_j = 0. \tag{47}$$

Note that each FOC defines implicitly a best reply function, and that such a function is negatively sloped, as was the case under Cournot competition with homogenous goods.

At the symmetric equilibrium, $q_i = q_j = q$, and by re-arranging the FOCs the equilibrium output is obtained:

$$q^C = \frac{\alpha}{2\beta + \gamma}. \tag{48}$$

By substitution one can then find the equilibrium output and profits, as:

$$p^C = \frac{\alpha\beta}{2\beta + \gamma}; \quad \pi^C = \beta(\frac{\alpha}{2\beta + \gamma})^2. \tag{49}$$

Price competition

If firms compete in prices, each firm $i = 1, 2$ chooses p_i so as to maximise its profits $\pi_i = q_i(p_i, p_j)p_i$ taking p_j as given, where $q_i(p_i, p_j) = a - bp_i + gp_j$. The FOCs are:

$$\frac{d\pi_i}{dp_i} = \alpha - 2bp_i + gp_j = 0. \tag{50}$$

The FOCs define the following two best reply functions in the plane (p_1, p_2):

$$R_1: \; p_2 = \frac{a - 2bp_1}{g}; \; R_2: \; p_2 = \frac{a + gp_1}{2b}. \tag{51}$$

Figure 8.5 illustrates the reaction functions and shows that they are positively sloped.[27] (Hence, the higher p_j the higher the price p_i which is the best reply to it.) It also illustrates the isoprofit curves of the two firms, that is the locus of the points such that a firm earns a given profit. For both firms, the further an isoprofit curve from the origin the higher the associated profits.[28]

INSERT Figure 8.5. Reaction functions under price competition

[27] Stability requires R_1 to be steeper than R_2, that is $2b/g > g/(2b)$, or $4b^2 > g^2$, always satisfied since $b > g$ by assumption.

[28] The isoprofit curves of firm 1 are described by the function $\pi_1 = (a - bp_1 + gp_2)p_1 = k$. By totally differentiating we obtain the slope of one such curve. First, take $d\pi_1 = (a - 2bp_1 + gp_2)dp_1 + gp_1 dp_2 = 0$, from which $dp_2/dp_1 = -(a - 2bp_1 + gp_2)/(gp_1)$, the isoprofit slope. (Note that along R_1 the numerator is zero, so crossing the reaction function the isoprofit must be flat.) Firm 2's isoprofits are given by $\pi_2 = (a - bp_2 + gp_1)p_2 = z$. Total differentiation gives $d\pi_2 = (a - 2bp_2 + gp_1)dp_2 + gp_2 dp_1 = 0$. Hence, one has $dp_2/dp_1 = -gp_2/(a - 2bp_2 + gp_1)$, the slope of firm 2's isoprofit. (Note that along the reaction function, R_2, the denominator is zero, so when crossing R_2 the isoprofit must be vertical.)

At the symmetric equilibrium, $p_i = p_j = p$, and by re-arranging the FOCs the equilibrium price is obtained:

$$p^B = \frac{a}{2b - g} = \frac{\alpha(\beta - \gamma)}{2\beta - \gamma}. \tag{52}$$

By substitution one can then find the equilibrium output and profits, as:

$$q^B = \frac{ab}{2b - g} = \frac{\alpha\beta(\beta - \gamma)}{2\beta - \gamma}; \quad \pi^B = \frac{a^2 b}{(2b - g)^2} = \frac{\alpha^2\beta(\beta - \gamma)}{(\beta + \gamma)(2\beta - \gamma)^2}. \tag{53}$$

It is worth noting that - although it was a good exercise to derive the equilibrium solutions - I could have found the solution of the price game just by noting the duality of the quantity and price problems, and substituting in the Cournot solutions already found above. Indeed, under quantity competition for firm i the problem is $\max_{q_i} \pi_i = (\alpha - \beta q_i - \gamma q_j) q_i$ given q_j, whereas under price competition the problem is $\max_{p_i} \pi_i = (a - bp_i + gp_j)p_i$, given p_j. These expressions are the dual of each other, as one can obtain the latter from the former by replacing q_i by p_i; α by a; β by b; γ by $-g$.[29]

Equilibrium comparison It is now possible to compare the results obtained under price and quantity competition. To do so, write the difference in the equilibrium prices as:

$$p^C - p^B = \frac{\alpha\gamma^2}{4\beta^2 - \gamma^2} > 0. \tag{54}$$

First, note that the difference is always positive: prices are always higher (and accordingly quantities are always lower) under quantity competition than under price competition, regardless of whether the goods are substitutes, independent, or complements.[30]

Second, $p^C - p^B$ increases with γ for given β: when $\gamma = 0$, the prices coincide (the markets being independent, it is irrelevant whether firms compete in quantities or in prices); when $\gamma = \beta$, the difference is highest, with $p^C - p^B = \alpha/3$ (the reader will have already recognised that this is the case corresponding to homogenous goods).

Strategic substitutes v. strategic complements We have seen above that under quantity competition the firms' reaction functions are negatively sloped, whereas under price competition they are positively sloped.[31] In what

[29] This also implies that Cournot competition with substitute (respectively complement) products is the dual of Bertrand competition with complement (respectively substitutes) products: $\gamma > 0$ corresponds to $g < 0$, and vice versa.

[30] One might also check another of the results obtained by Singh and Vives: consumer surplus and total surplus are (weakly) larger under price competition than under quantity competition. Profits are higher, equal or lower under quantity than under price competition, according to whether goods are substitutes, independent or complements.

[31] Note, however, that this holds for linear demand, but need not always be the case for more general demand functions.

follows, I briefly relate the slope of firms' reaction functions to a property of the profit function, and introduce the concepts of strategic substitutes and strategic complements, due to Bulow et al. (1985).[32]

Consider a simultaneous move game where each firm $i = 1, 2$ can choose actions a_i. Define $\pi_i(a_i, a_j)$ firm i's profit function, and assume that it is twice differentiable in both a_i and a_j, with $\partial^2\pi_i/(\partial a_i)^2 < 0$ on its domain to ensure a maximum exists.

As we know already, firm i's *reaction function* is given by the function $R_i(a_j)$ that identifies the best possible response a_i to any given action a_j of the rival:[33]

$$\frac{\partial \pi_i(R_i(a_j), a_j)}{\partial a_i} = 0. \tag{55}$$

Differentiation of (55) with respect to a_j gives:

$$\frac{\partial\left(\partial \pi_i/\partial R_i\right)}{\partial a_i}\frac{\partial R_i}{\partial a_j}da_j + \frac{\partial\left(\partial \pi_i/\partial a_i\right)}{\partial a_j}da_j = 0, \tag{56}$$

that can be re-written as:

$$\frac{\partial^2 \pi_i}{\partial a_i^2}\frac{\partial R_i}{\partial a_j} + \frac{\partial^2 \pi_i}{\partial a_i \partial a_j} = 0. \tag{57}$$

Denote $\partial R_i/\partial a_j \equiv R_i'$ as the slope of firm i's reaction function, and re-arrange it to give:

$$R_i' = -\frac{\partial^2 \pi_i/(\partial a_i \partial a_j)}{\partial^2 \pi_i/\partial a_i^2} = 0. \tag{58}$$

Therefore, the sign of the slope of the reaction function is given by the sign of $\partial^2\pi_i/(\partial a_i \partial a_j)$. (Recall that $\partial^2\pi_i/\partial a_i^2 < 0$.)

Bulow et al. (1985) introduce the following definition.

- *Actions are strategic substitutes if* $\frac{\partial^2 \pi_i}{\partial a_i \partial a_j} < 0$.
- *Actions are strategic complements if* $\frac{\partial^2 \pi_i}{\partial a_i \partial a_j} > 0$.

Actions are strategic substitutes (respectively, strategic complements) when an increase in a_j reduces (resp. raises) firm i's marginal profitability ($\partial\pi_i/\partial a_i$), thus leading firm i to choose a lower (resp. higher) a_i. This explains why the best reply functions are negatively (resp. positively) sloped.

As we shall see in section 4.1 below, the implications for an oligopoly model might turn out to be dramatically different according to whether the firms' actions are strategic complements or substitutes.

[32] This section follows closely Tirole (1988: 207-208).

[33] This best response is unique, due to the assumption that the profit function is strictly concave.

Stability of the equilibrium A property that is sometimes required of Nash equilibria in oligopoly models is that of *stability*. Stability of the equilibrium refers to the following question: would the equilibrium be reached through a sequence of movements along the reaction functions starting from any arbitrary point? Figure 8.6 illustrates the adjustment process involved by this thought experiment, under two different possibilities.

INSERT Figure 8.6. (a) Stable Nash Equilibrium; (b) Unstable Nash Equilibrium

Figure 8.6(a) illustrates a case of stable equilibrium (such as in the Cournot model analysed in this section). Suppose for instance that firm 1 produces a quantity q_1^0; firm 2's best reply would be to produce the quantity q_2^1; in turn, the best reply for firm 1 to the quantity q_2^1 would be to produce q_1^2; and so on, until the firms end up producing the quantities corresponding to the equilibrium pair at point E.

In Figure 8.6(b), instead, the equilibrium would not be stable, as movements along the reaction functions would push the firms further and further away from the equilibrium point E. Firm 2's best response to firm 1 producing quantity q_1^0 would be quantity q_2^1; in turn, the best reply for firm 1 to the quantity q_2^1 would be to produce q_1^2; however, firm 2's best reply would at this point be to produce nothing. The equilibrium point E cannot be reached through this adjustment process and is therefore said to be unstable.

As one can see from these two examples, the stability property is related to the slope of the reaction functions. In a game where firms 1 and 2 choose actions (a_1, a_2) and where reaction functions are drawn in the plane (a_1, a_2), for instance, it must be that R_1 is steeper than R_2. By recalling the expression of the slope of the reaction function, this condition can be written as $(\partial^2\pi_1/\partial a_1^2)(\partial^2\pi_2/\partial a_2^2) > (\partial^2\pi_1/\partial a_1\partial a_2)(\partial^2\pi_2/\partial a_2\partial a_1)$.

It is important to stress that the type of adjustment process that the stability concept supposes is very particular. The sequence of moves described above makes sense only if players are completely *myopic*, as each firm ignores the impact that its choice has on the rival's next move.[34]

Joint-profit maximisation If firms maximise joint profits and choose quantities, their problem is $\max_{q_1,q_2} \Pi$, where $\Pi = \pi_1 + \pi_2 = (\alpha - \beta q_1 - \gamma q_2)q_1 + (\alpha - \beta q_2 - \gamma q_1)q_2$. From the FOCs $d\Pi/dq_i = 0$ it follows $\alpha - 2\beta q_i - 2\gamma q_j = 0$. At the symmetric equilibrium, $q_1 = q_2$. Therefore, the FOCs simplify to:

$$q^M = \frac{\alpha}{2(\beta + \gamma)}. \quad (59)$$

By substitution one obtains equilibrium prices and per-firm profits as:

$$p^M = \frac{\alpha}{2}; \quad \pi^M = \frac{\alpha^2}{4(\beta + \gamma)}. \quad (60)$$

[34] Alternatively, one can think that each player does look ahead but attaches no importance to the future, that is it has a zero discount factor.

It can be checked that under price competition one gets exactly the same results:

$$p^M = \frac{a}{2(b-g)} = \frac{\alpha}{2}; \quad q^M = \frac{a}{2} = \frac{\alpha}{2(\beta+\gamma)}. \tag{61}$$

A differentiated goods model with some nice properties

The differentiated goods model used in the previous section has the advantage of being very easy to deal with. However, its extension from 2 to n firms (that I omit here for simplicity) is not completely satisfactory because aggregate demand increases with the number of firms (as well as with the degree of substitutability among products), at given prices. For this reason, I usually prefer to use another simple model, due to Shubik and Levitan (1980), where market size does not vary either with the degree of substitutability or the number of products. This is based on the following utility function for the differentiated products:[35]

$$U(q_1, .., q_i, ..q_n) = v\sum_{i=1}^{n} q_i - \frac{n}{2(1+\mu)}\left[\sum_{i=1}^{n} q_i^2 + \frac{\mu}{n}\left(\sum_{i=1}^{n} q_i\right)^2\right], \tag{62}$$

where q_i is the quantity of the i-th product, v is a positive parameter, n is the number of products in the industry, $\mu \in [0, \infty)$ represents the degree of substitutability between the n products.[36]

From the maximisation of the utility function subject to the income constraint, one can write $\partial U(\cdot)/\partial q_i - p_i = 0$, resulting in the inverse demand functions:

$$p_i = v - \frac{1}{1+\mu}\left(nq_i + \mu\sum_{j=1}^{n} q_j\right). \tag{63}$$

By inverting this system (see Appendix for details) we can find the following direct demand functions:

$$q_i = \frac{1}{n}\left[v - p_i(1+\mu) + \frac{\mu}{n}\sum_{j=1}^{n} p_j\right]. \tag{64}$$

Two nice properties of this demand function are: (1) aggregate demand $Q = \sum_{i=1}^{n} q_i$ does not depend on the degree of substitution among the products, as $Q = \sum_{i=1}^{n} q_i = v - \frac{1}{n}\sum_{i=1}^{n} p_i$. (2) In the case of symmetry $p_i = p_j = p$ aggregate demand does not change with the number of products n in the industry, as $Q = \sum_{i=1}^{n} q_i = v - p$.

[35] Of course, as above, consumers preferences can be expressed as $V = U(q_1, .., q_i, ..q_n) + y$, so that a partial equilibrium analysis is fully justified.

[36] A drawback of this utility function is instead that it does not allow to deal with both complements and substitutes.

Price competition Assume that all the firms have identical cost functions $C(q_i) = cq_i$, with $c < v$. Firm i's profits are given by $\pi_i = (p_i - c)q_i(p_1, .., p_i, .., p_n)$, where $q_i(\cdot)$ is given by (64). By writing the FOCs $\partial\pi_i/\partial p_i = 0$ and imposing symmetry ($p_i = p_j = p$ for all $j \neq i$) one obtains the "Bertrand" equilibrium prices as:

$$p_b = \frac{(v + c\,(n + n\mu - \mu))}{2n + n\mu - \mu}. \tag{65}$$

The quantity sold and the profits earned by each firm at equilibrium are found by substitution as:[37]

$$q_b = \frac{(v - c)\,(n + n\mu - \mu)}{n\,(2n + n\mu - \mu)}; \quad \pi_b = \frac{(v - c)^2\,(n + n\mu - \mu)}{(2n + n\mu - \mu)^2}. \tag{66}$$

Quantity competition

To study the case of quantity competition, write firm i's profits as $\pi_i = [p_i(q_1, .., q_i, .., q_n) - c]\,q_i$, where $p_i(\cdot)$ is given by (63). From the FOCs $\partial\pi_i/\partial q_i = 0$ and imposing symmetry ($q_i = q_j = q$ for all $j \neq i$) one obtains the "Cournot" equilibrium prices as:

$$q_c = \frac{(v - c)\,(1 + \mu)}{2n + n\mu + \mu}. \tag{67}$$

Prices and profits at the Cournot equilibrium are:

$$p_c = \frac{v(n + \mu) + cn\,(1 + \mu)}{2n + n\mu + \mu}; \quad \pi_c = \frac{(v - c)^2\,(1 + \mu)\,(n + \mu)}{(2n + n\mu + \mu)^2}. \tag{68}$$

The reader can check that - as in the previous model - quantity competition results in higher equilibrium prices and profits than price competition:

$$p_c - p_b = \frac{(v - c)(n - 1)\mu^2}{(2n + n\mu + \mu)(2n + n\mu - \mu)} \geq 0; \quad \pi_c - \pi_b = \frac{(v - c)^2\,(n - 1)^2\mu^3(2 + \mu)}{(2n + n\mu + \mu)^2(2n + n\mu - \mu)^2}, \tag{69}$$

the two equilibria coinciding when $\mu = 0$.

Joint-profit maximisation If firms maximise joint profits and choose prices, their problem is $\max_{p_1,.,p_i,.,p_n} \Pi = \sum \pi_i(p_1, ., p_i, ., p_n)$. From the FOCs $d\Pi/dp_i = 0$ and imposing symmetry one has:

$$p^M = \frac{v + c}{2}. \tag{70}$$

By substitution one obtains equilibrium prices and per-firm profits as:

$$q^M = \frac{v - c}{2}; \quad \pi^M = \frac{(v - c)^2}{4n}. \tag{71}$$

[37] Note that $\lim_{\mu\to\infty} p_b = c$ and $\lim_{\mu\to\infty} \pi_b = 0$: when goods become perfect substitutes, the equilibrium tends to the usual Bertrand case with homogenous goods.

Repeated product market interaction So far, we have analysed only one-shot games. This section looks at the case where firms still have only one strategic variable (be it price or quantity) but they repeatedly interact in the product market. In other words, I relax assumption A2 of the Bertrand game. It turns out that the result of the modified game is very different according to whether the game is repeated for a finite or infinite number of times.

3.2.2 Finite horizon

Consider the Bertrand game described in section 3.1.1, but with firms playing that basic game for $T+1$ periods, that is from period $t=0$ to period $t=T$, with T being a *finite* number.

Each firm $i=1,2$ wants to maximise the present discounted value of its profits, $\sum_{t=0}^{T}\delta^t\pi_{i,t}$, with δ being the discount factor[38] and $\pi_{i,t}$ the profit earned in period t. I am going to show the following result:

(Repeated Bertrand game, with finite horizon). If T is finite, the only subgame perfect Nash equilibrium of the game is the Bertrand equilibrium repeated T times.

Proof. This is a game of "almost perfect information" that can be solved by backward induction (see section 2.4.1). At the last period of the game, $t=T$, whatever happened in the previous periods, everything is as if the two firms were playing the one-shot Bertrand game. Therefore, the only equilibrium of the game will be the one-shot Bertrand price $p_{1,T}=p_{2,T}=c$.

At period $t=T-1$, players know that their current choices will not affect the equilibrium solution at the following period T. Therefore, whatever happened at periods $0,1,..,T-2$, the game they are playing at $T-1$ is effectively the same as if they were playing for the last time, and again the only equilibrium is the Bertrand equilibrium $p_{1,T-1}=p_{2,T-1}=c$.

The same reasoning can be applied to all previous periods, leading firms to choose $p_{1,t}=p_{2,t}=c$ at any period t. ■

This shows that when the Bertrand game is repeated a finite number of times, its outcome is exactly the same as the one-shot game, with firms setting prices at marginal cost and making no profits at any period. This result, however, holds only insofar as firms play a complete information game. If firms had *incomplete information* about their opponents (that is, if they were uncertain about their pay-offs), then the equilibrium outcome would not be the Bertrand equilibrium repeated T times. Although I am not going to analyse the Bertrand game under incomplete information, the interested reader can look at the incomplete information version of the chain-store paradox game in chapter 7, that shares similar features.[39]

[38] The discount factor measures the importance that a player attaches to the future: if $\delta=0$, then only current profits matter; if $\delta=1$, profits earned at any time in the future (however distant) hold the same importance as current profits. Since $\delta=1/(1+r)$, where r is the interest rate, $\delta=0$ corresponds to the case where $r\rightarrow\infty$, and $\delta=1$ to the case $r=0$.

[39] Kreps et al. (1982) have first formalised the repeated prisoner dilemma game (that is very similar to the Bertrand game) under incomplete information. See also Tirole (1988: 258-59).

A very different result arises also when firms repeatedly play the Bertrand game over an infinite horizon, as I show next.

3.2.3 Infinite horizon

Take now the number of times for which the price competition game is being played to be infinite.[40] Consider the following *trigger strategies*. Each firm sets the price p at the initial period $t = 0$. It sets price p at time t if both firms have set p in every period before t. Otherwise, each firm sets $p = c$ forever. In other words, each firm behaves in a "collusive" fashion as long as the rival does, but if one of them deviates from the "collusive" price, then the punishment is triggered and they both revert to the one-shot Bertrand equilibrium for the rest of the game.

These trigger strategies form an equilibrium if each firm's incentive compatibility constraint holds:[41]

$$\frac{\pi(p)}{2}(1 + \delta + \delta^2 + \delta^3 + ...) \geq \pi(p). \tag{72}$$

The LHS gives the present discounted value of the profits a firm receives if it colludes (i.e., if it follows the trigger strategy when the other firm does). In each period, the firm receives half of the aggregate monopoly profit. The RHS gives the present discounted value of the profits if the firm (optimally) deviates from the trigger strategy. By setting $p - \varepsilon$, all consumers will buy from the deviating firm, that will thus earn a profit $\pi(p - \varepsilon)$. For $\varepsilon \to 0$, this will therefore ensure profits very close to $\pi(p)$ in the period the deviation takes place. In the following period, however, the punishment occurs as both firms revert to the Nash equilibrium forever. Therefore, the deviating firm (as well as the rival) will make zero profit in all following periods of the game.

Note that $\sum_{t=0}^{\infty} \delta^t = 1/(1 - \delta)$. Hence, after simple algebra, the incentive constraint above becomes:

$$\delta \geq \frac{1}{2}, \tag{73}$$

that is, the price $p \in [c, p_m]$ can be sustained at the equilibrium provided that the discount factor is high enough.

There is little point in elaborating further here, since this is discussed at length in chapter 4. The main purpose of this section is just to show that prices above marginal costs might be sustained as the equilibrium of a game with infinite horizon (or uncertain final date). Note, however, that an important issue is that of the multiplicity of equilibria arising in this game. Any price

[40] Equivalently, suppose that the game has an uncertain final date, with the probability that the market exists in the following period being $\phi \in (0, 1)$. Call the firms' discount factor under this alternative interpretation d. One can then set $\delta = d\phi$, and carry out the analysis as in the text.

[41] The constraint is the same for both firms, due to their symmetry. See chapter 4 for collusion under asymmetries.

$p \in [c, p_m]$, that is between marginal cost and monopoly price, might be the equilibrium of the game (for high a enough discount factor). This leads to important policy implications that are discussed in chapter 4.

4 Oligopoly II: Dynamic games

So far, I have assumed that the only strategic variables available to firms were either prices or quantities, and that all their other characteristics were exogenously given. Clearly, this is not the case in reality, where firms also take a number of decisions that affect their costs and their products, and even the very choice of being in the market in the first place. When studying models where firms have several strategic variables, it is important to recognise that some are typically more "long-run" variables than others. Suppose for instance that we want to study a game where firms take entry, R&D investment and price decisions. It is reasonable to expect that price decisions are "short-run" decisions that can be revised relatively often, R&D investment decisions are more costly to revise, and entry decisions are probably "long-run" decisions that are taken only once and are taken as given forever afterwards. Accordingly, it would make sense to formalise the game played by the firms as a dynamic (or multi-stage) game where in the first stage firms choose whether to enter or not; in the second stage they decide how much R&D they want to carry out; and in the third stage what price they want to sell their products at.

Dynamic games have been described in general terms in section 2.4, and many such games are analysed throughout the book. For instance, chapter 2 analyses games where entry decisions are taken before (cost-reducing R&D or quality) investments and product market competition. Chapter 4 looks again at R&D-then-product market decisions when there are R&D spillovers. Chapter 5 analyses games where firms first decide whether to merge or not and then compete in the marketplace. In chapter 6 retailers make effort decisions or manufacturers take investment decisions before competition takes place; and in exclusive dealing models decisions are taken sequentially by an incumbent monopolist, buyers and a potential entrants. In chapter 7, again incumbents take some actions before potential entrants decide on entry, or new firms decide whether to continue operations or not.

In what follows, rather than offering yet other examples of dynamic games, I focus on the strategic effects that actions taken in the early stages of a game can have on oligopolistic interaction in the marketplace.

4.1 Strategic investments

In oligopoly models where firms take decisions sequentially, firms can act strategically, that is they can take actions in the early stages of the game that modify to their advantage some of the variables which are taken as given in the following stages. To understand how such strategic effects take place, consider the following game.[42]

In stage one of the game, Nature moves and picks the magnitude and sign of a "shock" parameter, $s \in \mathbb{R}$, that affects the marginal cost of firm 1, that equals $k - s$, in stage two.[43]

In stage two, firms 1 and 2 take product market decisions in a differentiated good market where consumers have utility function (43). Except for the shock that affects firm 1, firms would have the same marginal cost, k. However, to study possible entry-deterrence effects, assume that firm 2 also incurs a fixed cost f.

I consider two variants of the game: (i) firms choose quantities (in this case, decisions will be strategic substitutes) and (ii) firms choose prices (decisions will be strategic complements). I am interested in studying the impact of the shock s on the product market equilibrium (in particular, firms' profits) at stage two.

4.1.1 Quantity competition (strategic substitutes)

To analyse quantity competition, consider the inverse demands given by (44). Given the shock, the firms' problem is given by $\max_{q_1} \pi_1 = (\alpha - \beta q_1 - \gamma q_2 - k + s)q_1$ and $\max_{q_2} \pi_2 = (\alpha - \beta q_2 - \gamma q_1 - k)q_2 - f$. By taking the first derivatives, equaling them to zero and re-arranging, we can write the reaction functions in the plane (q_1, q_2) as follows:

$$R_1 : \; q_2 = \frac{\alpha - (k - s) - 2\beta q_1}{\gamma}; \; R_2 : \; q_2 = \frac{\alpha - k - \gamma q_1}{2\beta}. \tag{74}$$

Figure 8.7 illustrates the firms' reaction functions.

INSERT Figure 8.7. Effect of a shock ($s > 0$) that reduces firm 1's marginal cost: Strategic substitutes.

First, note that the reaction functions are negatively sloped: hence, in this model decisions are *strategic substitutes*. Then, note that the position of firm 1's reaction function is shifted by the shock s. In particular, if $s > 0$ as in the figure, R_1 shifts to the right to R_1'. The evaluate the strategic effect of the shock, it is convenient to distinguish between two cases. In the first, entry by firm 2 is not an issue (for instance, because $f = 0$ and the shock is such that at the intersection between R_1' and R_2 firm 2 sells a positive output). Call this *accommodation* case.[44] In the second case, firm 2's entry is an issue (for instance, because at the equilibrium absent the shock, firm 2 has positive *net* profits and a movement to a new equilibrium might put it on an iso-profit curve associated with negative profits, thus making entry unprofitable). Call this *entry deterrence* case. I analyse each case in turn.

[42] This is a very simplified, one-market-only version of Bulow et al. (1985). The other important reference for this section is Fudenberg and Tirole (1984).

[43] We shall see below that one could extend and re-interpret the model so that s is chosen by a firm, or by a third party such as a government.

[44] The term is drawn from Fudenberg and Tirole (1984), where the shock is endogenous. It refers then to an investment that will bring about a new equilibrium at which firm 2 has positive profits: in this sense, the investment accommodates entry.

Accommodation Assume that entry will be accommodated. A difficulty in assessing the impact of the shock which moves the equilibrium from E to E' is that it has two effects on firm 1's profits. First, the shock has a *direct effect* on firm 1's profits in that it decreases the costs of production of firm 1 at any pair of actions (q_1, q_2). Second, it has a *strategic effect* in that, since the shock is observed by its rival, it modifies the latter's choice. It is the second effect that we are interested in, as the first effect is independent of firm 2's response, and would take place even if firm 2 could not see the shock affecting firm 1.

To disentangle these two effects, reason in the following way. Once the shock occurs, it modifies firm 1's reaction function, shifting it to the right. The direct effect of the shock on firm 1's profits could therefore be seen as a movement which takes place absent any response on the side of firm 2. In other words, suppose that firm 2 does not change its action and still produces the quantity q_2^E. Then, firm 1 would end up in point D, that is on its new reaction function corresponding to the lower marginal costs, but where the quantity produced by firm 2 is unchanged. Firm 1's iso-profit passing by point D is associated with a profit level π_1^D.

However, firm 2 does observe the shock, and it knows that firm 1's reaction function is now R_1' and it will therefore revise its output accordingly, producing $q_2^{E'}$: firms end up in equilibrium E'. The strategic effect of the shock on firm 1's profits can therefore be seen as the move from D to E'. In words, the shock raises the output q_1 produced by firm 1. This decreases the marginal profitability of firm 2 (decisions being strategic substitutes), that will then lower its output q_2. In turn, lower q_2 increases firm 1's profits.

Since profits are higher (at E' firm 1 is on a lower iso-profit curve than at D), we can conclude that the strategic effect has a positive impact on π_1.

As for the strategic effect on firm 2's profitability, it is clearly negative. We have assumed no direct impact on firm 2's profits, and as the equilibrium shifts from E to E', firm 2 moves to a iso-profit function with lower profits.

Entry deterrence At the new equilibrium E', firm 2 is on an isoprofit function π_2' which corresponds to lower profits than the function π_2. Suppose that firm 2's *net* profits after the shock are $\pi_2' < 0$ (whereas they would be positive absent the shock: $\pi_2 > 0$). In this case, firm 1 will be certainly better off under the shock, since it would shift the duopoly equilibrium to a point where firm 2 would prefer to stay out. Hence, the shock would deter entry, and firm 1 would gain both because the shock has a positive direct effect on its profits (its costs are lower) and because it has the (positive) strategic effect to deter entry, and make it a monopolist.[45]

To sum up the quantity competition (strategic substitutes) case, the strategic effect of a cost-reducing shock on firm 1 is positive whether entry will be accommodated or deterred.

[45] We know from the accommodation case that fim 1 would be better off even if firm 2 stayed in the market. A fortiori, the shock will increase its profits when it makes firm 1 the only seller.

4.1.2 Price competition (strategic complements)

To analyse price competition, consider the system of direct demands given by (45). Firms' problems are $\max_{p_1} \pi_1 = (p_1 - k + s)(a - bp_1 + cp_2)$ and $\max_{p_2} \pi_2 = (p_2 - k)(a - bp_2 + cp_1) - f$. By taking the FOCs and re-arranging, we obtain the reaction functions in (p_1, p_2):

$$R_1 : p_2 = \frac{-a - b(k - s) + 2bp_1}{c}; \quad R_2 : p_2 = \frac{a + bk + cp_1}{2b}. \tag{75}$$

Figure 8.8 illustrates the firms' reaction functions.

INSERT Figure 8.8. The effect of a shock ($s > 0$) that reduces firm 1's marginal cost: Strategic complements

As we have already seen in section 3.2.1 above, the reaction functions are positively sloped: decisions are *strategic complements*. If $s > 0$ as in the figure, R_1 shifts to the left to R_1' following the cost-reducing shock.

Accommodation If entry is accommodated, the strategic effect associated with the shock might hurt firm 1. To see this, let me use the same trick as above to decompose direct and strategic effects on firm 1's profits. Absent any response to the shock from firm 2, that is if it set the same action p_2^E as in the case $s = 0$, firm 1 would be in point D. At such point, its profits would be π_1^D. But firm 2 modifies its price, determining a new equilibrium in E', where firm 1 has profits π_1'. The strategic effect is given by the move from D to E', and it is therefore negative for firm 1: $\pi_1' < \pi_1^D$ (any downward movement along the reaction function is associated with lower profits). What happens here is that the shock reduces firm 1's marginal cost and its price. Since price decisions are strategic complements in this model, the lower p_1 reduces firm 2's marginal profitability and accordingly its price p_2. In turn, the more aggressive pricing behaviour of firm 2 will reduce firm 1's profits.

The shock has also a negative effect on firm 2's profits, that are lower at E' than at E.

Entry deterrence Suppose that firm 2's *net* profits after the shock (that is, at E') are not only lower but also negative: $\pi_2' < 0$ (whereas they would be positive absent the shock: $\pi_2 > 0$). In this case, the shock will make firm 1 better off. This is because the shock would deter entry, and firm 1 would gain both because the shock has a positive direct effect on its profits (its costs are lower) and because it has the (positive) strategic effect to deter entry, and make it a monopolist.

To sum up the price competition (strategic complements) case, the strategic effect of a cost-reducing shock on firm 1 is negative when entry will be accommodated, but positive when it will be deterred.

Table 8.8 summarises the analysis.[46]

[46] Clearly, exactly the opposite results hold if the shock affected firm 1's negatively, that is if $s < 0$.

INSERT Table 8.8. Strategic effect of a shock that reduces firm 1's costs

Discussion and interpretation I have so far treated the shock as an exogenous variable, decided by nature or fate. However, the analysis above can be endogenised in a straightforward manner. Suppose for instance that s is determined by an investment decision of firm 1. For instance, a decision on R&D investment that reduces firm 1's costs and that is observable to firm 2: that is, a credible commitment to shift firm 1's reaction function.

The above analysis sheds light on whether oligopolistic interaction would push firm 1 to over- or under-invest, relative to a situation where the strategic effects were not taken into account and therefore would not affect the behaviour of firm 1's rival. Consider for instance the case where decisions are strategic substitutes: a reduction in its costs will have a positive strategic effect on firm 1, which will therefore be led to over-invest in cost-reducing activities (in this case, independently of whether entry of firm 2 is at stake or not).

Table 8.9, derived from the previous one, summarises the effects.

INSERT Table 8.9. Strategic investments to reduce firm 1's costs

Examples of strategic investments The strategic effects that I have illustrated above play an important role in a number of economic contexts analysed in this book.[47] For instance, manufacturers might use exclusive territorial clauses to induce retailers to choose higher prices, and in turn relax competition in the marketplace among both retailers and manufacturers when decisions are strategic complements (see chapter 6.3); an incumbent firm might over-invest in capacity, R&D and advertising to pre-empt entry by a new firm (see chapter 7.3); tie-in sales effectively make an incumbent more aggressive in the market where there exist competing sellers and can exclude the latter (chapter 7.3).

5 Appendix

I show here how to obtain the system of direct demands (64) from inverse demands (63).[48] The system of inverse demand functions can be written in matricial form as $p - v = -\frac{1}{1+\gamma} Aq$, where p, q are respectively the price and quantity $(n, 1)$ vector, v is a $(n, 1)$ vector having the scalar v in each entry, γ is a scalar, and A is a (n, n) matrix having element $n + \gamma$ on the diagonal and element γ everywhere off the diagonal. It is immediate to check that the direct demand functions can be rewritten in matricial form:

$$q = -(1 + \gamma)A^{-1} (p - v) . \tag{76}$$

Our problem is therefore to find A^{-1}, that is the inverse of matrix A.

[47] See Bulow et al. (1985) and Tirole (1988, pp. 328-336) for a number of other interesting applications.

[48] I am grateful to Felipe Cucker, great mathematician and gastronomer, who showed me a long time ago how to invert similar matrices.

Define $d = \frac{\gamma}{n}$. It is easy to check that $A = n\,(I + dO)$, where I is the identity matrix having 1 on the diagonal and 0 off the diagonal, and where O is the matrix with 1 in all its entries. Therefore, it must be $A^{-1} = \frac{1}{n}\,(I + dO)^{-1}$.

It turns out that $(I + dO)^{-1} = I - \left(\frac{d}{1+dn}\right) O$. We can check this by recalling that the product of a matrix by its inverse must be I. This requires a few steps, as follows.

$(I + dO)\left(I - \left(\frac{d}{1+dn}\right) O\right) = I + dO - \left(\frac{d}{1+dn}\right) O - \left(\frac{d^2}{1+dn}\right) O^2$.

One can immediately check that $O^2 = nO$. The previous expression can therefore be rewritten:

$I + dO - \left(\frac{d}{1+dn}\right) O - \left(\frac{nd^2}{1+dn}\right) O = I + \left[d - \frac{d}{1+dn} - \frac{nd^2}{1+dn}\right] O =$

$I + \left[\frac{d+d^2n}{1+dn} - \frac{d}{1+dn} - \frac{nd^2}{1+dn}\right] O = I + 0O = I$.

We can then conclude that the inverse of matrix A is given by:

$$A^{-1} = \frac{1}{n}\left[I - \left(\frac{\gamma}{n + n\gamma}\right) O\right]. \tag{77}$$

With few steps of algebra one can then simplify the expression $q = -(1 + \gamma)A^{-1}\,(p - v)$and check that it corresponds to the system of direct demand functions in the text.

Table 8.1. *A simple game*

B / A	b_1	b_2	b_3
a_1	2, 0	2, 5	1, 1
a_2	0, 2	0, 3	2, 2

Table 8.2. *The prisoners' dilemma game*

B / A	*High*	*Low*
High	10, 10	5, 15
Low	15, 5	6, 6

Table 8.3. *The battle of the sexes game*

B / A	*Indian*	*Thai*
Indian	3, 2	0, -1
Thai	-1, 0	2, 3

Table 8.4. *A pure coordination game*

A \ B	*Indian*	*Italian*
Indian	2, 2	0, 0
Italian	0, 0	1, 1

Table 8.5. *An asymmetric game*

B / A	p_L	p_H
p_L	0, 0	0, 0
p_H	0, -2	2, 0

Table 8.6. *The matching pennies game*

B / A	*Heads*	*Tails*
Heads	-1, 1	1, -1
Tails	1, -1	-1, 1

Table 8.7. *The entry deterrence game*

I / E	*Accommodate*	*Fight*
Enter	4, 5	-1, 0
Stay out	0, 10	0, 10

Table 8.8. *Strategic effect of a shock that reduces firm 1's costs*

Strategic substitutes	*Strategic complements*
Accommodation: $\pi_1 \uparrow$; $\pi_2 \downarrow$	Accommodation: $\pi_1 \downarrow$; $\pi_2 \downarrow$
Entry deterrence: $\pi_2 \downarrow$ ($\Rightarrow \pi_1 \uparrow$)	Entry deterrence: $\pi_2 \downarrow$ ($\Rightarrow \pi_1 \uparrow$)

Table 8.9. *Strategic investments to reduce firm 1's costs*

Strategic substitutes	*Strategic complements*
Accommodation: *over-invest*	Accommodation: *under-invest*
Entry deterrence: *over-invest*	Entry deterrence: *over-invest*

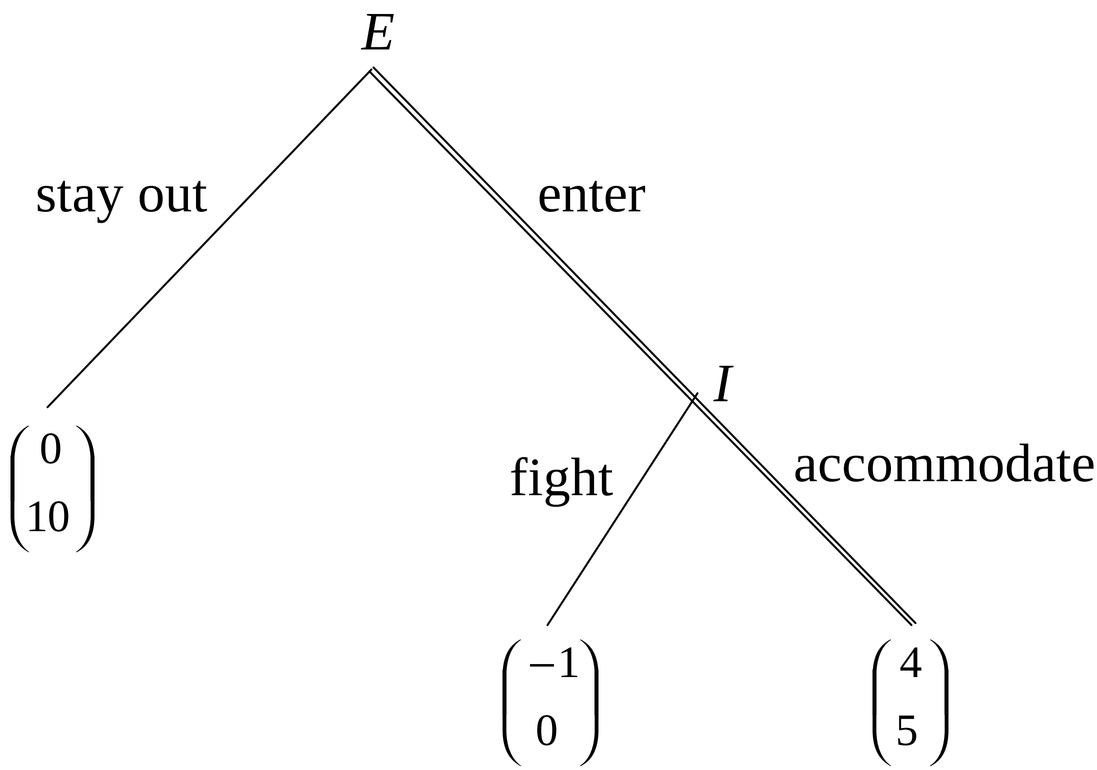

Figure 8.1. ***Extensive form of the entry deterrence game***

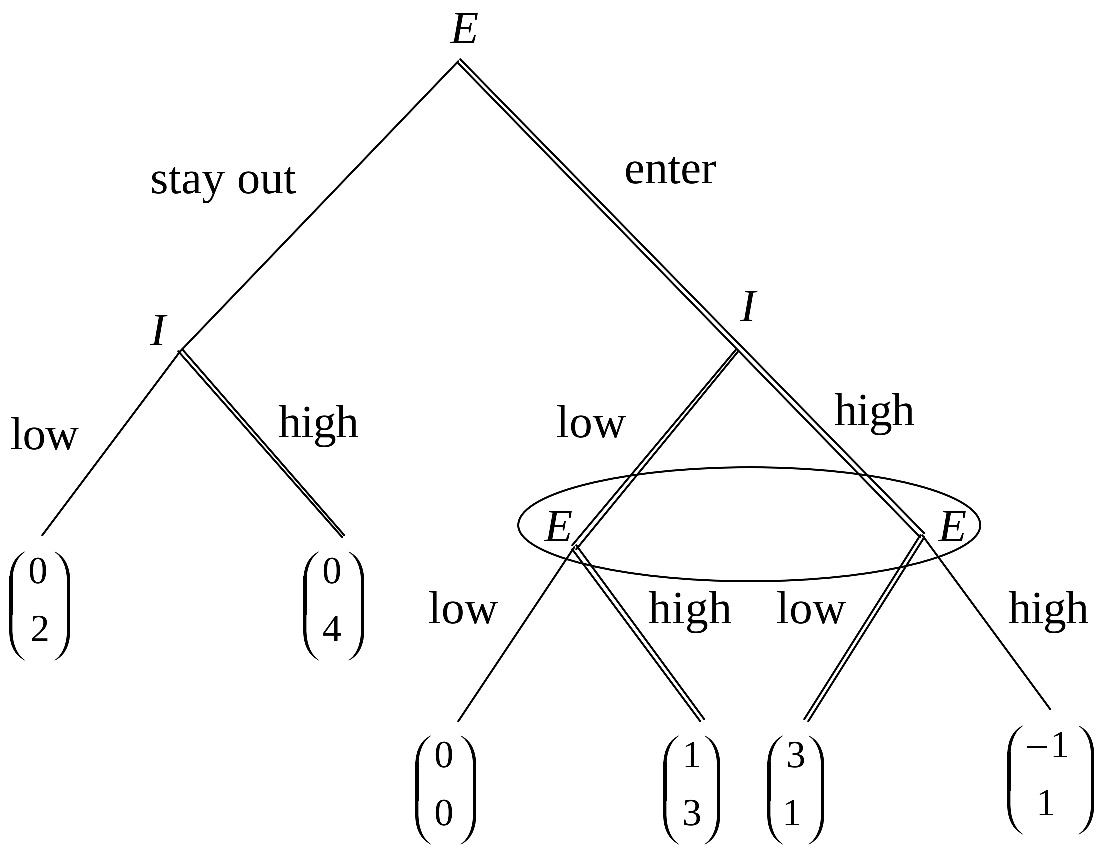

Figure 8.2. ***Extensive form of a quality game***

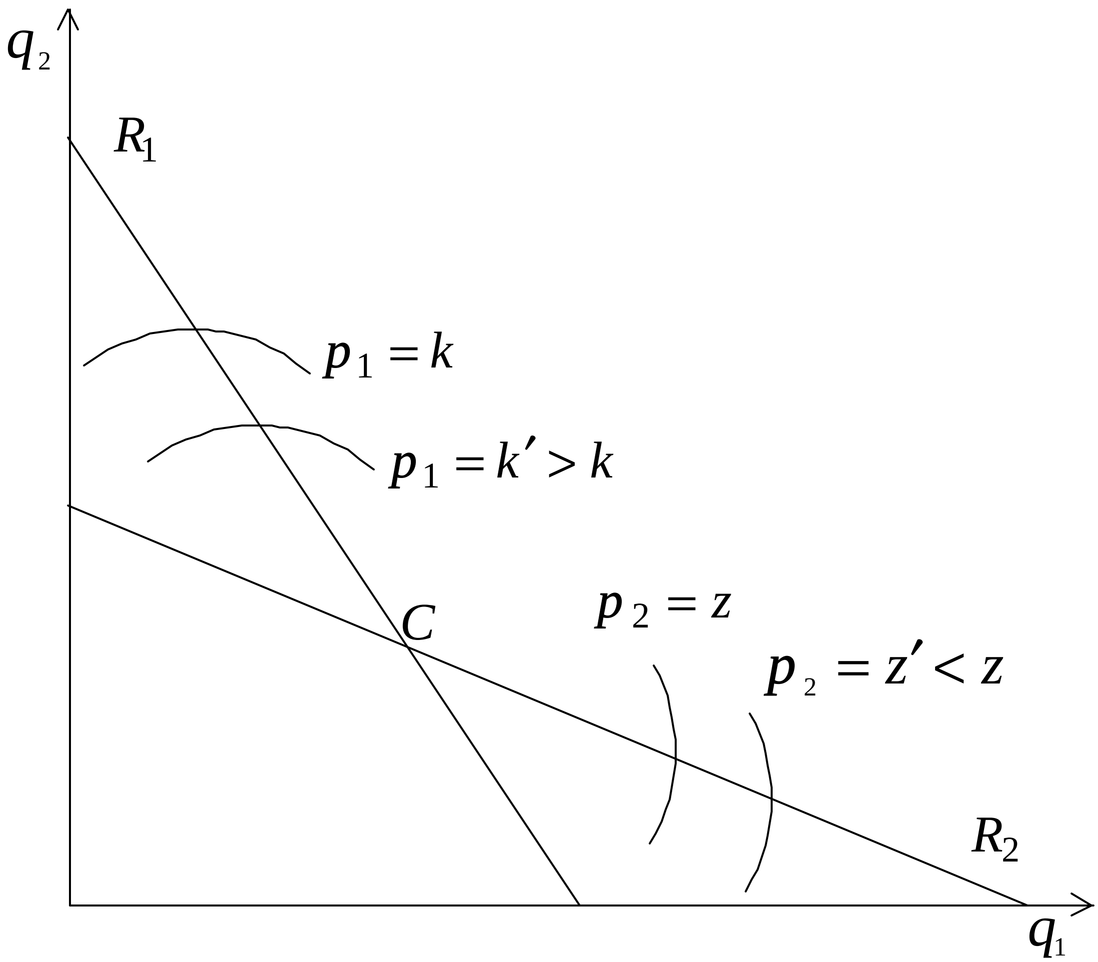

Figure 8.3. Reaction functions in the Cournot model

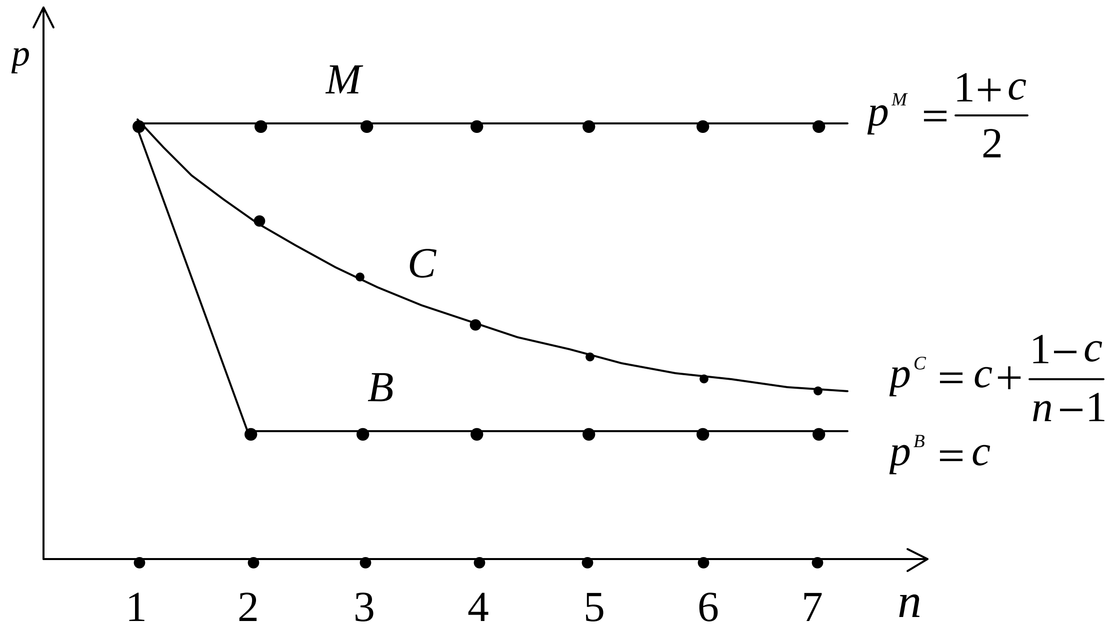

Figure 8.4. ***Toughness of product market competition: Bertrand (B), Cournot (C), and joint-profits maximization (M)***

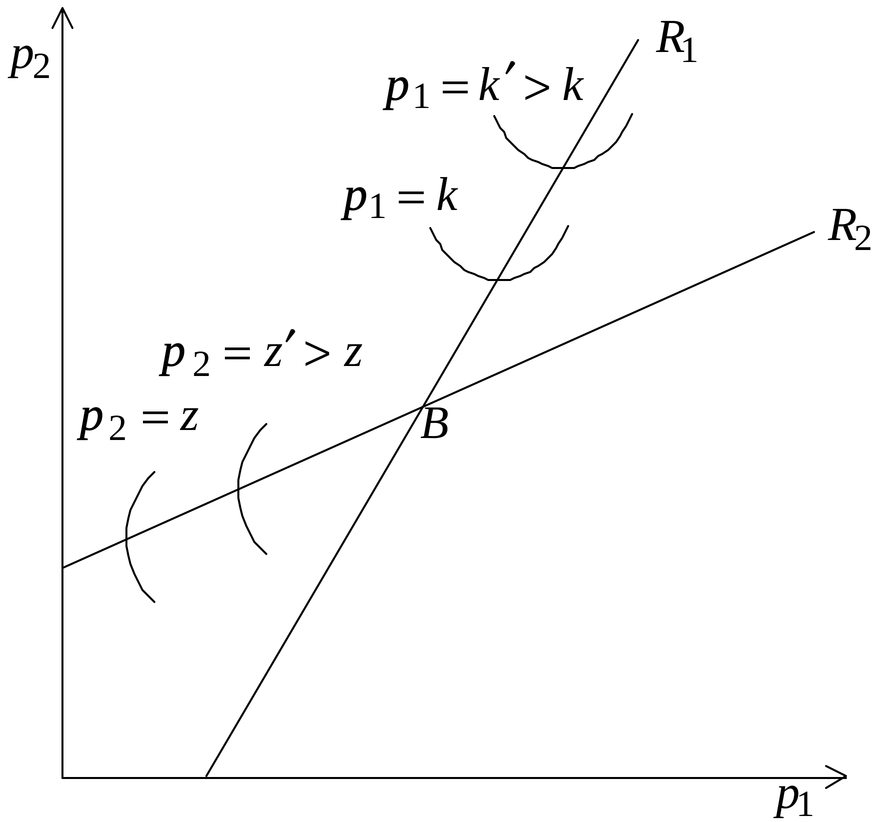

Figure 8.5. ***Reaction functions under price competition***

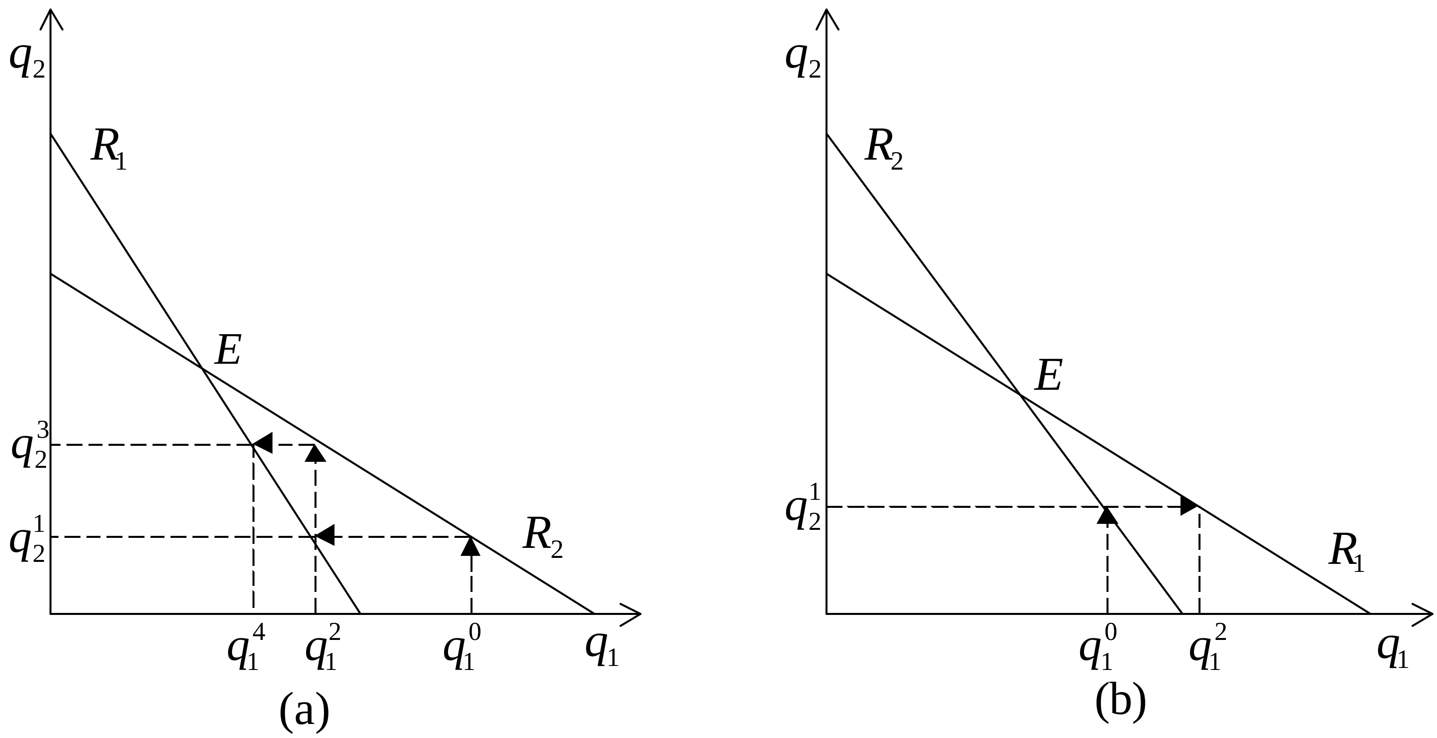

Figure 8.6. (a) ***Stable Nash Equilibrium; (b) Unstable Nash Equilibrium***

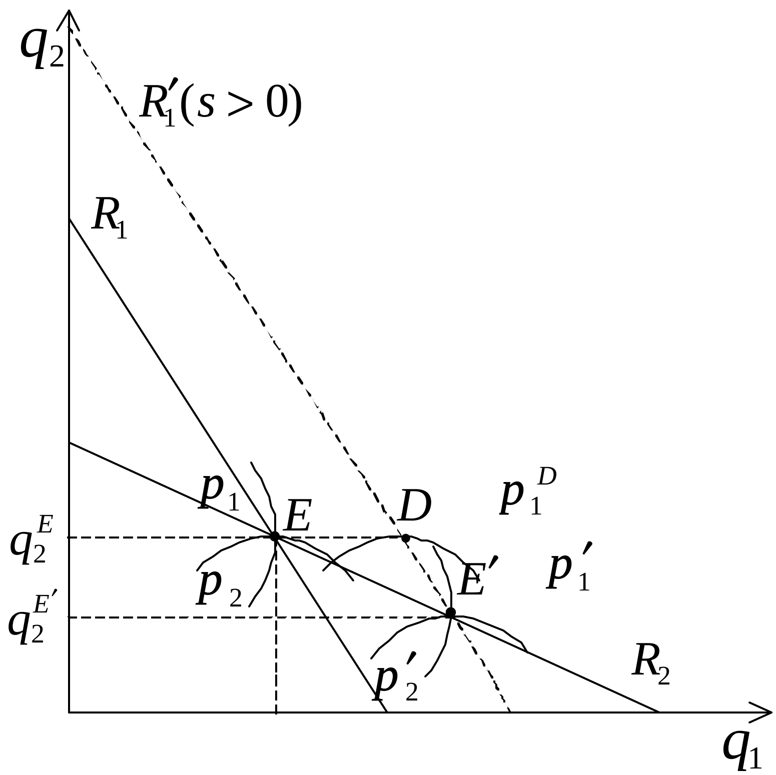

Figure 8.7. ***Effects of a shock (s > 0) that reduces firm 1's marginal cost: Strategic substitutes***

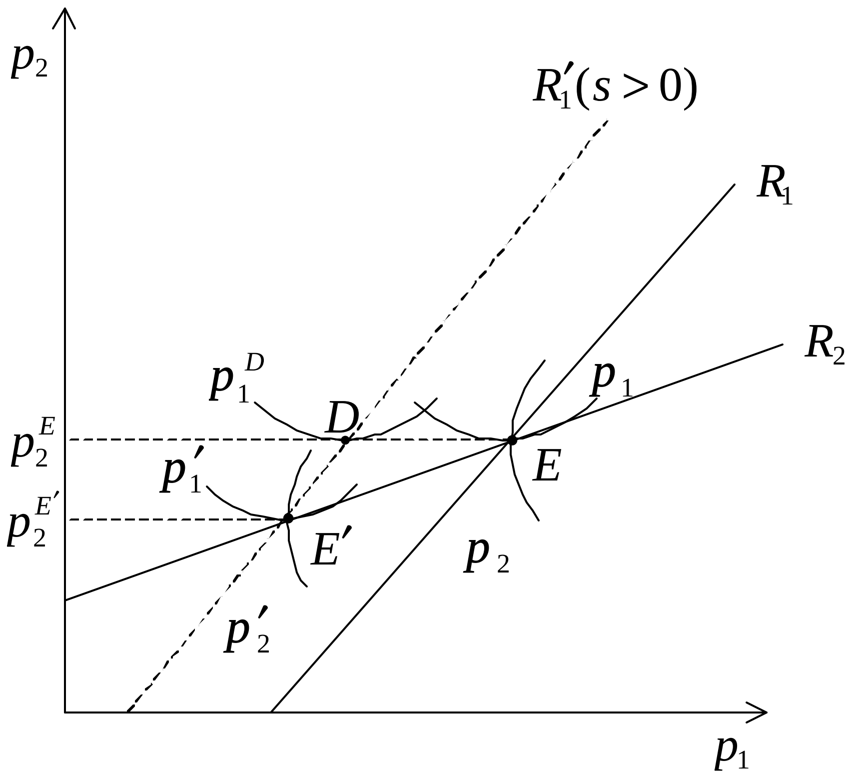

Figure 8.8. ***Effects of a shock (s > 0) that reduces firm 1's marginal cost: Strategic complements***