

Introduction to Industrial Organization

Supplemental Mathematical Sections

Luís Cabral

6 Perfect (and almost perfect) competition

A model of competitive selection

Section 6.3 presents a model of competitive selection by means of entry and exit. The main ideas of this model can be formalized in the following way. Suppose that each firm is characterized by a parameter θ . The average cost of a type- θ firm is given by $C(q_t)\theta$. That is, more efficient firms have a lower θ . In each period, the profit for a type- θ firm is given by

$$\Pi(q_t; \theta, t) = p_t q_t - C(q_t)(\theta + \epsilon_t), \quad (6.1)$$

where ϵ_t is a stochastic shock to the firm's productivity.¹ By assumption, ϵ_t has zero expected value and is independently distributed between periods and across firms.

At the beginning of each period, each firm decides whether to enter, remain inactive, exit, or remain active (depending on what its initial status is). Next, active firms decide how much to produce, before observing actual costs. Assuming that $\hat{\theta}_t$ is the firm's estimate of its θ , optimal output satisfies²

$$\max_{q_t} p_t q_t - C(q_t)\hat{\theta}_t. \quad (6.2)$$

By assumption, firms are price takers (the first assumption of the perfect competition model). Therefore, the first-order condition for profit maximization is given by

$$p_t = MC(q_t)\hat{\theta}_t,$$

¹The model of competitive selection is cast as one of heterogeneity in cost functions. However, it can be reinterpreted as one of product differentiation. Different values of θ would then correspond to different valuations by potential consumers. Dividing (6.1) by $\theta + \epsilon_t$, we get

$$\Pi(q_t; \theta, t) = \frac{p_t}{\theta + \epsilon_t} q_t - C(q_t).$$

This equation may be interpreted in the following way: consumers are willing to pay $\frac{p_t}{\theta}$ for the firm's product; in each period, the firm receives a signal of this valuation, namely $\frac{p_t}{\theta + \epsilon_t}$. The analysis of this model would be similar to the one we consider in the main text.

²Let x be a control variable and y a stochastic variable. Maximizing $E[f(x)g(y)]$ with respect to x is equivalent to maximizing $f(x)E[g(y)]$.

or

$$q_t^* = \Gamma(p_t/\hat{\theta}_t). \quad (6.3)$$

where q_t^* is optimal output and $\Gamma(\cdot)$ is the inverse of $MC(\cdot)$. Assuming that $MC(\cdot)$ is increasing, $\Gamma(\cdot)$ is increasing as well. We conclude that q_t is a decreasing function of the estimate for θ (that is, $\hat{\theta}_t$).

Consider the specific case when $C(\cdot)$ is a constant exponent function: $C(q) = q^\alpha$, where $\alpha > 1$. Then, $MC(q) = \alpha q^{\alpha-1}$, $\Gamma(p) = p^{\frac{1}{\alpha-1}}$, and, finally,

$$\Pi^*(q_t; \hat{\theta}, t) = \Omega \cdot \hat{\theta}_t^{-\frac{1}{\alpha-1}},$$

where Ω is an expression that does not depend on $\hat{\theta}_t$. This implies that not only output but also profits are decreasing in the value of $\hat{\theta}_t$.

7 Oligopoly

Comparative statics: exchange-rate devaluation

The first application considered in Section 7.5 corresponds to one of the firm's marginal cost increasing by 40%. Figure 1 summarizes the analysis in the main text and goes a bit further. The figure depicts the two levels of marginal cost, c and $1.4c$, as well as the increase in marginal cost, Δc . From the intersection of marginal revenue with marginal cost, we obtain the monopoly output level for each value of marginal cost. As we know, the monopoly output level gives the upmost value of the reaction curve. In the downward quadrant (fourth quadrant), we draw Firm 1's reaction curve (horizontal axis) as a function of Firm 2's output (downward vertical axis). This is basically Figure 7.10 but with a different orientation of the axes. By symmetry, the equilibrium is given by the intersection of the reaction curve and the -45° line. By drawing a line with slope 1 through the equilibrium point, we obtain, on the horizontal axis, the value of total output for each level of marginal cost. The difference is depicted with a horizontal arrow (pointing left) next to the horizontal axis. Finally, the difference in total output gives the difference in price, Δp , market with an arrow on the vertical axis.

The main points to remark in this figure are that: First, the increase in marginal cost implies a decrease in Cournot output that is greater than the decrease in monopoly output. (The decrease in monopoly output can be seen just to the left of the decrease in Cournot output.) Accordingly, the increase in monopoly price would be lower than the increase in Cournot equilibrium price. Second, under the new Cournot equilibrium, price increases by less than cost, as can be readily seen from the arrows along the vertical axis. Recall that, under perfect competition, price equals marginal cost, so that an increase in marginal cost implies an increase in price by the same amount.

To summarize, the price increase following an increase in marginal cost is greatest under perfect competition (or Bertrand oligopoly), lowest under monopoly, and intermediate under Cournot oligopoly. This is an important fact to which we return in Chapter 9.




Figure 1: Change in marginal cost, change in total quantity, and change in price.

8 Collusion

A model of secret price cuts

The following model formalizes the discussion in the first part of Section 8.2.³

In order to help us concentrating on the main issues, we assume a simple demand structure, namely an inelastic demand curve. Specifically, suppose that all consumers are willing to pay u for the (homogeneous) product sold by two duopolists. In each period, demand can be high (probability $1 - \alpha$) or low (probability α). When demand is high, $h = 1$ units can be sold at price u (or any lower price). When demand is low, only $l = 0 < h$ units can be sold. The probability that demand is high or low in each period is independent of what it was in the previous period. Moreover, *firms are unable to observe the state of market demand*; all they can observe is whether *their own* demand is high or low. Finally, for simplicity, assume that production costs are zero.

This set of assumptions about what firms can and cannot observe implies that equilibria of the sort considered before cannot be implemented. In fact, those equilibria assumed that firms could detect (with certainty) that rivals cheated from setting the agreed-upon price, whereas now we assume that firms cannot observe their rival's price decisions. However, by observing their own demand (and prices), firms can make inferences, imperfect as they might be, about the rival's past decisions. Is this sufficient to sustain collusion? The answer is, for

³This model is adapted from Tirole, Jean, *The Theory of Industrial Organization*, Cambridge, Mass: MIT Press, 1989, who in turn presents a simplification of the model proposed by Green and Porter, op. cit.

some parameter values, positive.

Consider the following equilibrium strategies. Firms start by setting $p = u$. If they receive a positive demand (namely, $\frac{1}{2}$), then they continue to set $p = u$, that is, they remain in the “co-operative phase”. If however one of the firms (or both) receives zero demand, then both firms enter into a “price war:” they set $p = 0$ during T periods and, after this period, revert to $p = u$ again (the co-operative phase).⁴

The condition for this to be an equilibrium is, as usual, that the expected payoff from playing the equilibrium strategy is greater than the payoff from deviating (by setting a slightly lower price and taking all of the market demand). Let V be expected equilibrium discounted profits starting in a period during the co-operative phase. Then

$$V = (1 - \alpha) \left(\frac{u}{2} + \delta V \right) + \alpha \delta^{T+1} V. \quad (8.1)$$

The first term on the right-hand side corresponds to the case when demand is high (probability $1 - \alpha$), whereas the second term corresponds to the case of low demand (probability α). If demand is high, then each firm receives current profits of $\frac{u}{2}$. Moreover, beginning next period, their continuation expected payoff is V , for there is no reason to start a price war. If, however, demand is low (probability α), then it is common knowledge that at least one of the firms receives zero demand, and that a price war will start in the next period. As a result, firms receive zero profits today (because demand is zero) and zero profits in the next T periods (because they engage in a price war). After these T periods, firms revert to the co-operative phase, so that their continuation expected payoff from then on is V .

If a firm deviates (during the co-operative phase), then its expected discounted payoff is

$$V' = (1 - \alpha)u + \delta^{T+1}V. \quad (8.2)$$

In words: if demand turns out to be high (probability $1 - \alpha$), then setting a slightly lower price gives the deviator a current profit of u (as opposed to $\frac{u}{2}$). However, regardless of what the state of demand is today, firms will certainly enter in a price war beginning in the next period. In fact, regardless of the state of demand, the rival firm (the non-deviator) receives zero demand today, the condition that triggers a price war. For this reason, expected future discounted payoff is simply $\delta^{T+1}V$.

The condition that the prescribed strategy constitutes an equilibrium is that $V \geq V'$. It can be shown that this condition simplifies into

$$1 \leq 2(1 - \alpha)\delta + (2\alpha - 1)\delta^{T+1}. \quad (8.3)$$

If $\alpha < \frac{1}{2}$, this condition is equivalent to $T \geq \underline{T}$, where \underline{T} is a positive number.

What is then the optimal equilibrium?⁵ In other words, what is the optimal value of T ? From (8.1), we can see that the equilibrium discounted payoff is decreasing in the value of T .

⁴Notice that, if a firm receives zero demand, then it is common knowledge that a price war is going to start, that is, it is common knowledge that one of the firms receives zero demand. In fact, either demand is low, in which case both firms receive zero demand, or one of the firms deviates from $p = u$, in which case the deviating firm knows that the rival receives zero demand.

⁵Optimality here is understood within the class of equilibria we are considering. It is possible to find collusive equilibria that perform better than the ones considered here.

This is intuitive, for the greater the value of T the longer the price wars will be; and firms don't like price wars. Therefore, the optimal value of T is the lowest value such that the equilibrium is stable, that is, the lowest value such that (8.3) holds, that is, $T = \underline{T}$.⁶

Demand fluctuations and collusion

The model below formalizes the discussion in the second part of Section 8.2.⁷

The model below is similar to the model of secret price cuts, with the difference that we now assume that in each period, before setting prices, firms observe the state of demand.⁸ We also make the simplifying assumption that $\alpha = \frac{1}{2}$, that is, the high- and the low-demand states are equally likely.

If the discount factor is sufficiently large, then the same kind of equilibrium as in the previous section is stable. In this equilibrium, firms set $p = u$ in every period, regardless of the state of demand. In fact, if a firm decides to go along with the agreed-upon equilibrium strategies, its expected discounted profits are

$$\frac{1}{2}ud + \frac{\delta}{1 - \delta} \frac{1}{2} \left(\frac{1}{2}uh + \frac{1}{2}ul \right),$$

where $d = h$ or $d = l$ depending on whether today's demand is high or low, respectively. Assuming that a deviation implies that firms switch to setting price equal to marginal cost, the payoff from deviation is simply ud , where, again, $d = h$ or $d = l$. The conditions that setting $p = u$ is an equilibrium are therefore

$$\begin{aligned} \frac{1}{2}uh + \frac{\delta}{1 - \delta} \frac{1}{2} \left(\frac{1}{2}uh + \frac{1}{2}ul \right) &> uh \\ \frac{1}{2}ul + \frac{\delta}{1 - \delta} \frac{1}{2} \left(\frac{1}{2}uh + \frac{1}{2}ul \right) &> ul, \end{aligned}$$

or simply

$$\delta > \frac{2}{3 + l/h} \equiv \bar{\delta}, \tag{8.4}$$

$$\delta > \frac{2}{3 + h/l} \equiv \underline{\delta}. \tag{8.5}$$

Notice that, since $h > l$, (8.4) implies (8.5). This is intuitive: the temptation to cheat on the agreement and set a slightly lower price is especially strong in periods of high demand. The condition for stability of the full-collusion agreement is therefore (8.4).

Suppose now that δ is lower than, but close to, $\bar{\delta}$. Clearly, full collusion cannot be an equilibrium. However, can *some* level collusion be sustainable? The answer is positive. The crucial point is that, if the price that firms agree to set is lower than the monopoly price (u), then the incentives for deviation are relatively smaller. In fact, suppose that firms decide to

⁶We are ignoring here the fact that T must be an integer.

⁷This model is adapted from Rotemberg, Julio, and Garth Saloner, "A Supergame-Theoretic Model of Price Wars During Booms," *American Economic Review* **76** (1986), 390–407.

⁸Firms can also observe past decisions by rival firms.

set price equal to $p_h < u$ during periods of high demand and $p_l = u$ during periods of low demand. The condition for no-deviation during periods of high demand is now given by

$$\frac{1}{2}p_h h + \frac{\delta}{1-\delta} \frac{1}{2} \left(\frac{1}{2}p_h h + \frac{1}{2}ul \right) \geq p_h h$$

or simply

$$p_h < \frac{\delta l/h}{2-3\delta} u.$$

Since firms would like to set prices as high as possible, we get

$$p_h = \frac{\delta l/h}{2-3\delta} u. \quad (8.6)$$

Substituting $\bar{\delta}$ for δ in (8.6), we get, as expected, $p = u$. Any value $\delta < \bar{\delta}$ yields $p_h < u$. Assuming that δ is close to $\bar{\delta}$, p_h will in turn be close to u . This implies that the constraint (8.5) is still satisfied. We thus conclude that setting $p = p_h$ in periods of high demand and $p = u$ in periods of low demand is an equilibrium.

9 Market structure and market power

Concentration and market power, I

In Section 9.1, it is shown, by means of two graphics, that market power decreases as the number of firms increases. We now prove this fact algebraically. Recall from Chapter 7 that Firm 1's reaction function is given by

$$q_1^*(q_2) = \frac{a-c}{2b} - \frac{q_2}{2}.$$

More generally,

$$q_1^*(q_2) = \frac{a-c}{2b} - \frac{Q_{-1}}{2}.$$

In a symmetric equilibrium, $q_1 = q^N$ and $Q_{-1} = (n-1)q^N$, where n is the number of firms. We thus have

$$q^N = \frac{a-c}{2b} - \frac{(n-1)q^N}{2}.$$

Solving for q^N , we get

$$q^N = \frac{a-c}{(n+1)b}.$$

It can be easily seen that $q^N \rightarrow 0$ as $n \rightarrow \infty$. Moreover, if Q^N is total equilibrium output, we have

$$Q^N = nq^N = \frac{n}{n+1} \cdot \frac{a-c}{b}. \quad (9.1)$$

As $n \rightarrow \infty$, we get

$$Q^N \rightarrow \frac{a-c}{b} = Q^C,$$

that is, total output converges to the perfect competition level.

Derivation of Table 9.1

This is the table that gives efficiency loss as a function of the number of firms.

From the previous mathematical section, we know that

$$Q^N = \frac{n}{n+1} \cdot \frac{a-c}{2b}.$$

Substituting in the demand curve, $p = a - bQ$, we get

$$\begin{aligned} p^N &= a - b \left(\frac{n}{n+1} \cdot \frac{a-c}{2b} \right) \\ &= \frac{a+nc}{n+1}. \end{aligned}$$

(Note that, as expected, $p^N \rightarrow c$ as $n \rightarrow \infty$.)




Figure 2: Allocative efficiency loss in Cournot equilibrium.

The efficiency loss is given by area A in Figure 2. Analytically, this is given by

$$\begin{aligned} A(n) &= \frac{1}{2}(p^N - p^C)(Q^C - Q^N) \\ &= \frac{1}{2} \left(\frac{a+nc}{n+1} - c \right) \left(\frac{a-c}{2b} - \frac{n}{n+1} \cdot \frac{a-c}{2b} \right) \\ &= \frac{1}{2} \cdot \frac{a-c}{n+1} \cdot \frac{a-c}{2b(n+1)} \\ &= \frac{1}{2} \cdot \frac{(a-c)^2}{2b(n+1)^2}. \end{aligned}$$

Efficiency loss under monopoly is simply given by

$$A(1) = \frac{1}{2} \cdot \frac{(a-c)^2}{2b(1+1)^2}.$$

Therefore

$$\frac{A(n)}{A(1)} = \frac{(1+1)^2}{(n+1)^2}.$$

Substituting for different values of n , we obtain the values in Table 9.1.

Derivation of Equation 9.1

Equation (9.1) can be derived as follows. Firm i 's profit is given by

$$\pi_i(q_1, \dots, q_n) = P(Q)q_i - C_i.$$

The first-order condition for profit maximization is

$$P'q_i + p - MC_i = 0,$$

(where P' is the derivative of $P(Q)$ with respect to Q), or simply

$$p - MC_i = -P'q_i.$$

Dividing both sides by p , and rearranging, we get

$$\begin{aligned} \frac{p - MC_i}{p} &= \frac{-P'q_i}{p} \\ &= \frac{-P'Q}{p} \cdot \frac{q_i}{Q} \\ &= \frac{s_i}{\epsilon}, \end{aligned}$$

where $s_i \equiv q_i/Q$ is firm i 's market share and $\epsilon \equiv (P'Q/p)^{-1}$ is the price elasticity of demand.

Finally,

$$\begin{aligned} L &\equiv \sum_{i=1}^n s_i \frac{p - MC_i}{p} \\ &= \sum_{i=1}^n s_i \frac{s_i}{\epsilon} \\ &= \frac{\sum_{i=1}^n s_i^2}{\epsilon} \\ &= \frac{H}{\epsilon}. \end{aligned}$$

as claimed in the text.

10 Price discrimination

General pricing schedules and two-part tariffs

Section 10.3 starts with the derivation of an optimal two-part tariff (when all consumers are of the same type). There, we *assume* that the seller sets a two-part tariff, and then derive the optimal two-part tariff. However, it can be shown that a two-part tariff is an optimal pricing schedule within the broader set of non-linear pricing schedules.

Suppose that each consumer is willing to pay $W(q)$ in order to consume q units.⁹ Suppose the seller sets a pricing schedule $T(q)$, which in principle can be any function. A particular case corresponds to **linear pricing**, $T(q) = pq$, the case we have considered in previous chapters. If the seller is not restricted to linear pricing, what is the optimal $T(q)$?

Since $W(q)$ is the willingness to pay, the net utility for each consumer is given by $W(q) - T(q)$. This net utility must be non-negative, or else consumers would not buy at all. In fact, from the seller's point of view, $T(q)$ should be such that net utility is exactly zero. If that were not the case, then the seller could increase the fixed part of $T(q)$ by a little without losing customers, thus increasing total profits. We thus have $W(q) - T(q) = 0$, or simply $W(q) = T(q)$. But then the seller's total profit, $T(q) - C(q)$, can be rewritten as $W(q) - C(q)$, which is simply total surplus.

The above implies that, *in order to maximize profits, the seller should choose a pricing schedule that: (a) induces the consumers to buy the quantity that maximizes total surplus; (b) leaves each consumer with zero net utility.*

There are many functions $T(q)$ which imply that consumers choose the efficient q (that is, the one that maximizes total surplus) and receive zero net utility. The simplest one is undoubtedly the two-part tariff $T(q) = CS(c) + cq$, as derived above.

Two-part tariff with multiple consumer types

The first part of Section 10.3 is very simplistic in that it assumes all consumers have the same demand. Let us now consider the more realistic assumption that there are two different types of consumers.¹⁰ In the example of a telecommunications operator, this would correspond to a group of consumers (group 1) that make very few phone calls and a group of consumers that make many phone calls (group 2). Figure 3 depicts the demand curves of each type of consumer.

If all consumers were of type 1, then the optimal two-part tariff would consist of $p = c$ and $f = CS(c)$, as we saw before. Can this be an optimal tariff when there are two types of consumers, type 1 and type 2, in equal number? The answer is "no." Consider the alternative two-part tariff $p = p_1$, $f = CS_1(p_1)$, where $CS_1(p)$ is the consumer surplus of type-one consumer when price is p . Let us compare this solution with the previous solution $p = c$, $f = CS_1(c)$. If the consumer is of type 1, then the new solution implies a loss of $A + B$ in terms of fixed fee but a gain of A in terms of variable revenues (for now price is greater than marginal cost); this corresponds to a net loss of B . If, instead, the consumer is of type 2, then the new solution

⁹Normally, consumer preferences are characterized by the demand curve $D(p)$, where p is price. But since we are now dealing with more general pricing schemes, it is better to start from the concept of willingness to pay.

¹⁰This is still not entirely realistic: in the real world, there are as many demand curves as there are consumers. However, from a qualitative point of view, the analysis of the two-type case is equivalent to the analysis of the n -type case.




Figure 3: Two-part tariff with two types of consumers.

implies, again, a loss of $A + B$, but now a gain of $A + B + C$, or a net gain of C . Since C is greater than B and there is an equal number of consumers of each type, it follows that the two-part tariff $p = p_1$, $f = CS_1(p_1)$ implies a greater profit than $p = c$, $f = CS_1(c)$. It can also be shown that this solution is better than setting $p = c$, $f = CS_2(c)$ (where $CS_2(c)$ is consumer surplus for a type-2 consumer), a two-part tariff which would have the implication that only type-2 consumers make buy.

What if there are two sellers competing with each other? The profit for seller 1, for example, is given by $\pi_1(T_1, T_2) = T_1 q_1(T_1, T_2)$, where $q_1(T_1, T_2)$ is seller 1's demand, which in turn is a function of both sellers' pricing schedules. This problem can be restated by taking into account that, from the consumer's perspective, what matters is the net utility it can get from the pricing schedule T_i . Let this net utility be $u_i = (T_i)$ in the case of firm i . We thus have $\pi_1 = T_1 q_1(u_1, u_2)$.

This restatement of Firm 1's problem suggests a simple two-step procedure for its optimization problem: first, for each value of u_1 , determine the pricing schedule T_1 that maximizes profits; and second, determine the level of net utility u_1 to give consumers (as a function of the level of net utility u_2 offered by the rival firm). This restatement is important because it shows that, even with competition, the above results remains true that Firm 1's optimal pricing schedule is one that maximizes total surplus for its consumers. In fact, maximizing $T_1 - C_1$ subject to the constraint that $W_1 - T_1 = u_1$ is equivalent to maximizing $W_1 - C_1 - u_1$ with respect to T_1 , which in turn implies an efficient solution, as shown above.

Damaged goods

The model below formalizes the analysis of damaged goods, presented in Section 10.4.

Suppose there are two types of buyers. High-valuation buyers, of which there is a fraction α , are willing to pay \bar{v}_h for the "good" version of the product and \underline{v}_h for the "damaged" version. Low-valuation buyers are willing to pay \bar{v}_l and \underline{v}_l , respectively. We assume that high-valuation buyers are willing to pay much more for the "good" version than low-valuation buyers, but approximately the same for the "damaged" version. We also assume that all valuations are greater than cost, which, in turn, is assumed to be zero. To summarize,

$$\bar{v}_h \gg \bar{v}_l > \underline{v}_h \approx \underline{v}_l > c = 0.$$

Let us first consider the case when no damaged good is sold. The seller has two options: either it sets $p = \bar{v}_h$ or its sets $p = \bar{v}_l$. The first price leads to profits of $\alpha\bar{v}_h$, whereas the second yields \bar{v}_l . Given our assumption that $\bar{v}_h \gg \bar{v}_l$, we conclude that the optimal price without price discrimination is $p = \bar{v}_h$, and profits $\alpha\bar{v}_h$.

Let us now consider the case when the seller introduces a damaged version of the product. Naturally, the seller wants to induce the low-valuation buyers to purchase the low-quality version and the high-valuation buters to purchase the high-quality one. Let the price of the low-quality version be $\underline{p} = \underline{v}_l$. This is the maximum price the seller can possibly charge the low-valuation buyers. What is the maximum price the seller can charge the high valuation buyers? A high-valuation buyers will be indifferent between the two products if $\bar{v}_h - \bar{p} = \underline{v}_h - \underline{p}$. Therefore, the maximum price the seller can charge for the good version of the product is $\bar{p} = \bar{v}_h - (\underline{v}_h - \underline{p}) = \bar{v}_h - (\underline{v}_h - \underline{v}_l)$.

Given these prices, $\bar{p} = \bar{v}_h - (\underline{v}_h - \underline{v}_l)$ and $\underline{p} = \underline{v}_l$, profits are given by $\alpha\bar{p} + (1 - \alpha)\underline{p} = \alpha[\bar{v}_h - (\underline{v}_h - \underline{v}_l)] + (1 - \alpha)\underline{v}_l$.

From our assumption that $\underline{v}_h \approx \underline{v}_l$, it follows that total profit is greater under price discrimination than under uniform pricing. What is particularly interesting in this example is that all buyers, as well as the seller, are (weakly) better off with the introduction of a damaged good. In fact, low valuation buyers receive zero surplus under price discrimination (they pay a price equal to their valuation), the same as under no price discrimination (in which case they do not make any purchase). High-valuation buyers are strictly better off: they pay $\bar{p} = \bar{v}_h - (\underline{v}_h - \underline{v}_l)$ under price discrimination, less than what \bar{v}_h , the price paid under no price discrimination.

Durable goods

The model below formalizes the analysis of durable goods, presented in Section 10.4.

Suppose there are two types of buyers, in equal number (one million of each, say). High-valuation buyers are willing to pay up to v_h for one unit of a durable good. Low-valuation buyers are only willing to pay v_l , where $0 < v_l < v_h$. Both buyers discount the future according to the factor $\delta < 1$: one dollar tomorrow is worth δ dollars today. The seller has a production cost of zero and needs to decide on prices in period 1 and 2.

Suppose the seller is to set the same price in both periods. The alternatives are to set $p = v_h$, yilding a profit of v_h (that is, v_h millions); or to set $p = v_l$, yielding a profit of $2v_l$. Suppose that $v_h > 2v_l$. It follows that, if a single price is to be set, then the seller's optimum is $p = v_h$.

Suppose now that we are in period 2 and that all high-valuation buyers have made a purchase in the first period. Only low-valuation buyers are left in the market. Since v_l is greater than cost, it is optimal for the seller to lower price to v_l for an additional profit of v_l . In total, the seller's profits are then $v_h + v_l$, which is greater than the profit that it would get if it were not to price discriminate.

The problem with this solution is that high-valuation buyers will see through it and wait for the second period. Given that second-period price is v_l , the maximum the seller can charge high-valuation buyers in the first period is given by $v_h - p_1 = \delta(v_h - p_2)$, the condition that a high-valuation buyer is indifferent between buying in the first period and buying in the second period. This can be solved to

$$p_1 = (1 - \delta)v_h + \delta p_2 = (1 - \delta)v_h + \delta v_l.$$

Total profits are then given by $(1 - \delta)v_h + \delta v_l + v_l$. If the value of δ is sufficiently close to 1, that is, if buyers are sufficiently patient, then the above expression is approximately equal to $2v_l$, which in turn is lower than monopoly profits under no price discrimination.

11 Vertical relations

Two part tariff and efficiency

In Section 11.1, we show that the double-marginalization problem can be solved if the upstream firm can set a two-part tariff. Let us restate this problem in a more formal way. The manufacturer's profit is given by

$$\pi_M = (w - c)D(p) + f.$$

The manufacturer will set as a franchise fee the maximum that the retailer is willing to pay. This is given by the retailer's profit:

$$f = (p - w)D(p).$$

Substituting for f in the previous equation we get

$$\pi_M = (w - c)D(p) + (p - w)D(p) = (p - c)D(p) = \pi_V(p).$$

That is, the manufacturer's profit is effectively the profit of a vertically integrated firm, $\pi_V(p)$. The manufacturer should therefore set the w that induces the retailer to set the price that maximizes $\pi_V(p)$. This corresponds to $w = c$.

Retailer competition and optimal wholesale price

Section 11.2 in the text briefly explores the effect of retailer competition. Here we restate the same ideas in a more formal way.

In order to consider intermediate levels of competition, we introduce a slightly more general formulation than in the text above, one where the product sold by the two retailers is not homogeneous. Let the demand directed to retailer i be given by $D(p_i, p_j)$. The assumption that products are substitutes is that D_2 , the derivative of D with respect to the second argument, is positive: if my rival increases its price, then my demand increases. The case of homogeneous product, considered above, corresponds to the limit when D_2 is equal to D_1 in absolute value, where D_1 is the derivative of D with respect to the first argument.¹¹ By contrast, the case of little substitutability corresponds to D_2 being close to zero.

As seen before, if fixed fees are possible, then the upstream firm's problem consists of maximizing total industry profits. These are given by

$$\pi_V = D(p_1, p_2)(p_1 - c) + D(p_2, p_1)(p_2 - c),$$

whereas R_i 's profits are

$$\pi_i = D(p_i, p_j)(p_i - w_i) - f_i,$$

¹¹Note however that D_2 is positive, whereas D_1 is negative.

where w_i and f_i are the wholesale price and the fixed fee paid by R_i , respectively. The first-order condition for R_i 's profit maximization is

$$\frac{\partial D(p_i, p_j)}{\partial p_i}(p_i - w) + D(p_i, p_j) = 0.$$

In a symmetric equilibrium, this simplifies to

$$D_1(p - w) + D = 0. \tag{11.1}$$

Likewise, in a symmetric equilibrium, the first-order condition for the maximization of total profits is

$$D_1(p - c) + D + D_2(p - c) = 0. \tag{11.2}$$

Comparing (11.1) with (11.2), we see that, at $w = c$, the latter includes the extra term $D_2(p - c)$. This implies that, at the point where (11.2) is satisfied ($p = p^M$), the left-hand side of (11.1) is negative. This in turn implies that, starting from p^M and $w = c$, an individual firm would have an incentive to reduce price. In order to counteract that incentive, M should set $w > c$, to the point where

$$D_1(p^M - w) + D = 0,$$

so that each downstream firm will want to set price at the monopoly level.

In this formulation, the degree of competition is proportional to the degree of substitutability across products. Specifically, the greater D_2 , the greater the degree of competition (the more a change in the price set by R_i (that is, p_i) will affect the demand for R_j 's product (that is, $D(p_j, p_i)$). Comparing (11.1) and (11.2), we see that a higher D_2 implies a greater gap between R_i 's first-order condition and M 's first-order condition. As a result, when D_2 is higher, w must be higher in order to induce the downstream firms to set $p = p^M$. We thus conclude that the greater the degree of retail competition, the greater the optimal wholesale price.

Investment externalities and indirect control

The model below formalizes the ideas discussed in Sections 11.3 and 11.4.

Let demand for retailer R_i be given by $D(p_i, p_j, s_i, s_j)$. s_i is an investment by R_i that increases demand (advertising, sales service, etc.). We assume that both $D_3 > 0$ and $D_4 > 0$, that is, R_i 's sales effort increases both its demand and that of its rival.¹²

R_i 's profit is given by $\pi_i = D(p_i, p_j, s_i, s_j)(p_i - w_i) - s_i - f_i$. In a symmetric equilibrium, the first-order condition for the maximization of π_i with respect to s_i is

$$D_3(p - w) = 1. \tag{11.3}$$

Total industry profits, in turn, are given by $D(p_i, p_j, s_i, s_j)(p_i - c) + D(p_j, p_i, s_j, s_i)(p_j - c) - s_i - s_j$. In a symmetric equilibrium, the first-order condition for the maximization of total profits with respect to s_i is

$$(D_3 + D_4)(p - c) = 1. \tag{11.4}$$

¹² D_i denotes the derivative of D with respect to the i th argument.

Suppose that $D_2 = 0$, that is, there is no price competition between downstream firms. The optimal wholesale price, in terms of inducing optimal pricing by downstream firms, is then $w = c$. Comparison of (11.3) to (11.4) shows that, if $w = c$, then there is a difference between the firm-level and the industry-level first-order conditions. At the privately optimal level of s (given by (11.3) the left-hand side of (11.4) is greater than one: a higher level of s would be desirable.

The model above also formalizes the ideas presented in Section 11.4. Consider again the case when demand for R_i is given by $D(p_i, p_j, s_i, s_j)$. Suppose now that $D_4 = 0$, so that there are no externalities with respect to the level of s_i . The first-order condition for the maximization of π_i with respect to p_i will look similar to (11.1), whereas the corresponding condition for total profits will be like (11.2). If price competition is very intense, that is, if $D_1 \approx -D_2$, realignment of private and collective incentives implies that w be set at a level close to p^M . But then the left-hand side of (11.3) will be very small, in particular, smaller than the left-hand side of (11.4) *even if* $D_4 = 0$.

Contract renegotiation

The analysis we have developed so far is subject to a crucial caveat: we have implicitly assumed that the upstream firm can commit to a fixed set of contracts with each downstream firm. Suppose, however, that this is not the case. In particular, suppose that the manufacturer initially agrees with the retailers on the optimal contract (f^*, w^*) . Moreover, assume that the degree of retail competition is not as extreme as Bertrand competition, so that the optimal wholesale price is below monopoly price and the fixed fee is strictly positive.

What will happen if the manufacturer can secretly offer each retailer R_i the possibility of renegotiating its contract? The manufacturer could then approach retailer R_1 and argue along the following lines: Now that M and R_2 have agreed on a contract with a positive fixed fee f^* and a wholesale price w^* close to p^M (the same as M and R_1 initially agreed), there are mutual gains from M and R_1 renegotiating their agreement. The new deal would correspond to a lower wholesale price and a higher fixed fee. Under this new deal R_1 increases profit and so does M . How is it possible that both M and R_1 become better off if the initial set of contracts maximizes total profits? Clearly, it must be the case that R_2 becomes worse off, and in fact such is the case: by paying a lower wholesale price, R_1 will decrease price and steal market share from R_2 . While this reduces R_2 's profit, the fixed fee it has to pay is still f^* , so M does not suffer from R_2 's profit reduction.

But would R_1 accept such an offer? In fact, one thing that R_1 might conjecture is that after renegotiating this contract and promising to pay M a higher fixed fee, M will then approach R_2 and attempt to renegotiate that contract as well, in which case the renegotiated deal proposed by M is not so attractive to R_1 . This type of reasoning will only stop at an equilibrium in which downstream firms know that it is not in the interest of M to renegotiate any of the deals. It can be shown that this equilibrium corresponds to a wholesale price lower than w^* , and, consequently, a lower profit level for M .¹³

In order for a set of contracts to be an equilibrium, it must be that there is no mutual gain

¹³Hart, Oliver, and Jean Tirole, "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity* (Microeconomics), 205–276, 1990.

from renegotiation for the manufacturer, M , and each retailer R_i . M 's profits are given by

$$\pi_M = D(p_i, p_j)(w_i - c) + D(p_j, p_i)(w_j - c) + f_i + f_j.$$

R_i 's profits, in turn, are given by

$$\pi_i = D(p_i, p_j)(p_i - w_i) - f_i,$$

and so

$$\pi_M + \pi_i = D(p_i, p_j)(p_i - c) + D(p_j, p_i)(w_j - c) + f_j.$$

The condition that there is no room for M and R_i to jointly improve their profits (given the contract signed between M and R_j) is that the derivative of $\pi_M + \pi_i$ with respect to p_i be zero. Notice that, strictly speaking, what M and R_i would be negotiating is the value of w_i , not the value of p_i . But since different values of w_i lead to different values of p_i , we may as well assume that M and R_i agree directly on the value of p_i . In a symmetric equilibrium ($p_i = p_j = p$, $w_i = w_j = w$), the condition that $\partial(\pi_M + \pi_i)/\partial p_i = 0$ becomes

$$D_1(p - c) + D + D_2(w - c) = 0. \tag{11.5}$$

Notice the contrast between (11.5) and (11.2). Since $w < p$, the third term in (11.5) is lower than the third term in (11.2). Since, by definition, (11.2) is satisfied for $p = p^M$, the left-hand side of (11.5) is negative for $p = p^M$: M and R_i would increase their joint profits by decreasing p_i . In order for (11.5) to be satisfied, it must be $p < p^M$ (so that the first term is less negative than with $p = p^M$).

In the above “renegotiation-proof” equilibrium, the upstream firm is victim of a sort of “curse”: The reason why M offers the downstream firms the possibility to renegotiate contracts is that by doing so it can increase profits. But, as a result of this, the upstream firm ends up making less profits than it would by simply offering f^* and w^* . In other words, if the upstream firm could commit not to renegotiate its original contracts, then it would be better off.¹⁴

One solution to the manufacturer’s commitment problem is given by **Resale Price Maintenance (RPM)**, whereby the manufacturer imposes a minimum price on retailers. Under RPM, the incentives for renegotiation cease to exist. In fact, as we saw in the previous section, renegotiation would imply setting a lower wholesale price, a higher fixed fee, and (implicitly) *a lower retail price*.¹⁵ But with RPM such new contract would not work.¹⁶

Why can the manufacturer commit to RPM and not to a set of contracts with retailers? The reason lies in that a retail price is more easily observable than a contract between manufacturer and retailer. In other words, the manufacturer could include in its contract with retailer R_1 the

¹⁴Notice the analogy between the manufacturer’s problem and the problem faced by a durable good’s monopolist (considered in the previous chapter): in both cases lack of commitment not to renegotiate terms of sale results in a worse outcome.

¹⁵We say “implicitly” because the two-part contract does not directly specify a retail price. However, a retail price reduction is required if the contract is to benefit the retailer.

¹⁶This idea was put forward by O’Brien, Daniel P, and Greg Shaffer, “Vertical Control With Bilateral Contracts,” *Rand Journal of Economics* **23** (1992), 299–308.

clause that “no retailer will set a price below p^M .” If retailer R_2 violates this clause, R_1 will in principle be able to observe it. However, observing a change in f_2 and w_2 , the fixed fee and the wholesale price paid by R_2 to the manufacturer, is more problematic—the scope for secret contract renegotiation is greater than the scope for evading an industry-wide “regulation.”

RPM is not the only way of solving the upstream firm’s commitment problem. An alternative is to award **exclusive territories**. This is a vertical restraint whereby each retailer is allocated a given territory that other retailers have no access to. For example, car manufacturers have an exclusive dealer in each European country. The German *Fiat* dealer, for instance, is not allowed to sell cars in France.

An exclusive territory solves the commitment problem. It allows the manufacturer to go back to the previous solution ($f = \pi^m$, $w = c$) and suggest that each retailer R_i set $p = p^M$. This would indeed be the retailer’s optimal choice, for it knows that, by contractual agreement, there won’t be any competition at the retail level: the retailer is, effectively, a monopolist with marginal cost $w = c$, and thus setting the monopoly price p^M is indeed its optimal choice.¹⁷

To summarize: If the manufacturer is unable to commit to fixed contractual terms with each retailer, then the manufacturer’s market power is dissipated in contract renegotiation and downstream competition. Vertical restraints such as RPM or exclusive territories allow the manufacturer to recover its market power.

12 Product differentiation

Hotelling equilibrium

Section 12.2 presents the Hotelling model in a graphical way. We now derive the same results algebraically.

We begin with a derivation of Firm 1’s demand curve. For generic prices (p_1, p_2) , the indifferent consumer’s location x is given by

$$p_1 + tx = p_2 + t(1 - x)$$

or

$$x = \frac{1}{2} + \frac{p_2 - p_1}{2t}$$

As seen in the text, Firm 1’s demand is given by all consumers to the left of the consumer located at x , that is

$$d_1 = x = \frac{1}{2} + \frac{p_2 - p_1}{2t}$$

(As expected, $p_1 = p_2$ implies $d_1 = \frac{1}{2}$.)

Firm 1’s profit is given by

$$\pi_1 = (p_1 - c) \left(\frac{1}{2} + \frac{p_2 - p_1}{2t} \right).$$

¹⁷As in the case of double marginalization, vertical integration would also solve the problem. In fact, vertical integration always solves the incentive problems in downstream pricing. Vertical integration does not take place more frequently because there are also negative incentive problems resulting from large, integrated firms. See Chapter 3.

The first-order condition for profit maximization, $\frac{\partial \pi_1}{\partial p_1} = 0$, implies

$$\frac{1}{2} + \frac{p_2 - p_1}{2t} - (p_1 - c)\frac{1}{2t} = 0.$$

Solving for p_1 , we get

$$p_1 = \frac{c + t}{2} + \frac{p_2}{2},$$

which is Firm 1's reaction function. At a symmetric equilibrium, $p_1 = p_2 = p^N$, where p^N is given by

$$p^N = \frac{c + t}{2} + \frac{p^N}{2},$$

or simply

$$p^N = c + t.$$

Product positioning

Below, we formalize the main ideas in Section 12.3.

The direct and the strategic effects can be derived algebraically as follows. Let $(p_1(l_1, l_2), p_2(l_1, l_2))$ be the equilibrium of the second stage (pricing) game. Notice that equilibrium prices depend on locations. Substituting these equilibrium prices on Firm 1's profit function, we get Firm 1's profit as a function of locations:

$$\pi_1 = \pi_1 \left(l_1, l_2, p_1(l_1, l_2), p_2(l_1, l_2) \right).$$

How does Firm 1's profit vary when l_1 varies? Notice that an increase in l_1 means moving closer to Firm 2 (given our assumption regarding relative locations).

$$\begin{aligned} \frac{d\pi_1}{dl_1} &= \frac{\partial \pi_1}{\partial l_1} + \frac{\partial \pi_1}{\partial p_1} \frac{\partial p_1}{\partial l_1} + \frac{\partial \pi_1}{\partial p_2} \frac{\partial p_2}{\partial l_1} \\ &= \frac{\partial \pi_1}{\partial l_1} + \frac{\partial \pi_1}{\partial p_2} \frac{\partial p_2}{\partial l_1}. \end{aligned}$$

In equilibrium, Firm 1 chooses an optimal price, and so it must be that $\partial \pi_1 / \partial p_1$ is zero. This implies that the second term in the first equality is zero, leaving us with two terms. These two terms correspond to the direct and the strategic effects, respectively. It can be shown that the first term is positive. In fact, we have done so based on Figure 12.4. The first part of the second term is positive: $\partial \pi_1 / \partial p_2 > 0$. This is quite general and intuitive: given that the products are substitutes, Firm 1 gains when Firm 2 increases its price. Finally, the second part of the second term is normally negative: $\partial p_2 / \partial l_1 < 0$. The idea is that locating closer to Firm 2 makes the latter price more aggressively, resulting in a lower equilibrium price. Together, these two inequalities combine into a negative strategic effect.

13 Advertising

Advertising as a signal

The model below formalizes the discussion in the second part of Section 13.1.

Consider the following simple model of advertising as a signal.¹⁸ A firm launches a new product. From the consumers' perspective, the product can be of high quality or low quality (with equal probability). Consumers do not know the firm's product quality until after buying the product. Moreover, consumers are willing to pay \bar{v} for a high-quality product and \underline{v} for a low-quality product. Finally, suppose that production cost is given by c and that

$$\bar{v} > c > \frac{1}{4}\underline{v} + \frac{3}{4}\bar{v}. \quad (13.1)$$

Because of the second inequality, if consumers have no information about the firm's quality, then no sale will take place. In fact, the expected willingness to pay (the maximum price a firm will be able to charge in the first period) is $(\underline{v} + \bar{v})/2$. If the firm's quality is high, then even by charging this price in the first period it would only get a total revenue of $(\underline{v} + \bar{v})/2 + \bar{v}$ (after the first period, consumers learn that the firm's product is of high quality, and are thus willing to pay \bar{v} in the second period). But (13.1) implies that this revenue is not sufficient to cover total production costs, $2c$. The same reasoning would apply (even more strongly) in the case when the firm's quality is low.

Now enter advertising. Suppose that, if the firm's quality is high, then it spends $a^* = \bar{v} - c$ on advertising. If, instead, the firm's quality is low, then it spends zero on advertising. Given this strategy, consumers should conclude that advertising means the firm's product is of high quality; accordingly, consumers should be willing to pay \bar{v} . Zero advertising, by contrast, implies that the firm's product is of low quality; accordingly, consumers should be willing to pay \underline{v} .

Can this be an equilibrium? In other words, is the firm choosing an optimal strategy? If the firm's product is of high quality, then, in equilibrium, it receives a profit of $\bar{v} - c - a^*$ in the first period and $\bar{v} - c$ in the second period. Since $a^* = \bar{v} - c$, this adds up to a total of $\bar{v} - c$, which is positive (first inequality in (13.1)). If the firm were to choose a different level of advertising, including $a = 0$, then consumers would believe the firm to be of low quality.¹⁹ The maximum price it would be able to sell is \underline{v} . Its total profit would thus be $\underline{v} - c + \bar{v} - c$, which is lower than $\bar{v} - c$, the total profit it gets by setting $a = a^*$.

Consider now the case when the firm's product is of low quality. In equilibrium it gets zero profit: if $a = 0$, then no price above cost would convince consumers to buy the product. The alternative is to spend $a = a^*$, thus (erroneously) convincing consumers that the product is of high quality. Under this scenario, the firm would be able to sell for \bar{v} in the first period. Eventually, consumers would learn that the product is of low quality, which implies that

¹⁸This model is a simplified version of Kihlstrom, Richard E, and Michael H Riordan, "Advertising as a Signal," *Journal of Political Economy* **92** (1984), 417–450, who in turn extend seminal work by R. Nelson, op. cit.). See also Migrom and Roberts, op. cit.

¹⁹Since, in equilibrium, only the levels $a = 0$ and $a = a^*$ are observed, there is no particular consumer "belief" that needs to be imposed when the firm chooses a different value of a . We assume that consumers believe the firm's product to be of low quality.

second period profits are zero. Total profit would thus be given by first period profit, that is, $\bar{v} - c - a^*$, which is equal to zero. We conclude that the firm can do no better than the equilibrium strategy, both when its quality is high and when its quality is low: the equilibrium strategy is indeed an equilibrium strategy.

The Dorfman-Steiner formula

In what follows, we derive Equation 13.1. Assuming constant marginal cost c , the firm's profit function is given by

$$(p - c)D(p, a) - a,$$

where D , the demand curve, is a function of price and advertising expenditures (in \$). The first-order conditions for profit maximization are given by $\partial\pi/\partial p = 0$ and $\partial\pi/\partial a = 0$:

$$q + (p - c) \frac{\partial D}{\partial p} = 0$$

$$(p - c) \frac{\partial D}{\partial a} - 1 = 0.$$

These equations can be transformed into

$$\frac{1}{p}(p - c) \frac{\partial D}{\partial p} = -\frac{q}{p}$$

$$\left(\frac{a}{pq}\right) (p - c) \frac{\partial D}{\partial a} = \frac{a}{pq},$$

or simply

$$\frac{p - c}{p} = \frac{1}{-\frac{\partial D}{\partial p} \frac{p}{q}} = \frac{1}{\epsilon}$$

$$\frac{a}{R} \equiv \frac{a}{pq} = \left(\frac{p - c}{p}\right) \left(\frac{\partial D}{\partial a} \frac{a}{q}\right).$$

Together, the above equations imply (13.1).

14 Entry costs, market structure, and welfare

Derivation of Equation 14.1

Let the inverse demand curve be given by

$$P = a - Q/S.$$

S is a measure of market size. If the only difference between market A and market B is that the latter is twice the size of the first, then the demand curve in market B is as in A but with a value of S that is twice that of market A 's.

Each firm's profit, in case it enters the market, is given by

$$\Pi = Pq_i - F - cq_i = (a - Q/S - c)q_i - F.$$

The first-order condition for profit maximization is

$$a - Q/S - c - q_i/S = 0$$

In a symmetric equilibrium, we have $q_i = q = Q/n$, where n is the number of active firms. It follows that

$$a - nq/S - q/S = 0,$$

or

$$q = \frac{a - c}{n + 1}S.$$

Equilibrium price is given by

$$\begin{aligned} p &= a - nq/S \\ &= a - n \frac{a - c}{n + 1}. \end{aligned}$$

Substituting in the profit function, we get

$$\begin{aligned} \Pi(n) &= (p - c)q - F \\ &= \left(a - n \frac{a - c}{n + 1} - c \right) S \frac{a - c}{n + 1} - F \\ &= \frac{a - c}{n + 1} S \frac{a - c}{n + 1} - F \\ &= S \left(\frac{a - c}{n + 1} \right)^2 - F. \end{aligned}$$

Free entry and welfare

The model below formalizes the discussion in Section 14.3. Suppose that the inverse demand function is given by $P(Q)$ and that each firm's cost function is $C(q)$ (this generalizes the case considered in the text). Finally, assume that firms decide sequentially whether or not to enter the industry.

The question we are trying to address is the relation between the optimal number of firms, n^* , and the equilibrium number of firms, \hat{n} .

The optimal number of firms maximizes total welfare, which is given by

$$W(n) \equiv \int_0^{nq_n} P(x)dx - nC(q_n) - nF, \quad (14.1)$$

where, as before, n is the number of firms and q_n each firm's output (given that there are n active firms).

The effect of additional entry on welfare is given by

$$W'(n) = P(nq_n) \left(n \frac{\partial q_n}{\partial n} + q_n \right) - C(q_n) - nMC(q_n) \frac{\partial q_n}{\partial n}.$$

The equilibrium number of firms, \hat{n} , is given by the zero-profit condition

$$P(\hat{n}q_{\hat{n}})q_{\hat{n}} - C(q_{\hat{n}}) = 0.$$

Substituting this in (14.1), we get

$$W'(\hat{n}) = \hat{n}(P(\hat{n}q_{\hat{n}}) - MC(q_{\hat{n}})) \frac{\partial q_n}{\partial n}. \tag{14.2}$$

We conclude that, if margins are positive ($P > MC$) and there is a business stealing effect ($\partial q_n / \partial n < 0$), then, at the equilibrium level of entry, further entry would reduce social welfare; and, conversely, less entry would increase welfare. In other words, there is excessive entry in equilibrium.

Notice that, by using differential calculus, this result abstracts from the fact that n must be an integer value. When this problem is taken into consideration, examples may be found where there is insufficient entry in equilibrium. For example, suppose that marginal cost is constant and that duopoly price is equal to marginal cost (Bertrand competition). In this case, even if the entry cost is very small, the equilibrium number of firms is just one. However, if the entry cost is indeed very small, then society would be better off with a second competitor.

15 Strategic behavior, entry and exit

Deterrence by capacity expansion

Section 15.1 develops the idea of entry deterrence by capacity expansion. Below we present a more detailed, formal analysis of that problem

Assume first that Firm 2's cost of entry is very small. Then Firm 1 knows that no matter what output level it sets (within reasonable limits), Firm 2 will enter.²⁰ What is, then, Firm 1's optimal output?

Firm 1 knows that Firm 2 will choose an optimal output *given Firm 1's preset output level*. Specifically, Firm 1 anticipates that Firm 2 will choose $q_2 = q_2^*(q_1)$, where $q_2^*(q_1)$ is Firm 2's reaction function as in the Cournot model. In fact, as we saw in Chapter 7, $q_2^*(q_1)$ gives Firm 2's optimal output given Firm 1's output. Figure 4 depicts this function. It also depicts some of Firm 1's **isoprofit curves** in the map (q_1, q_2) . Firm 1's π_1' isoprofit curve, for example, includes all pairs (q_1, q_2) such that Firm 1's profits are equal to π_1' , and likewise for π_1'' and π_1''' . Notice that Firm 1's profits are decreasing in its rivals output, so $\pi_1' > \pi_1'' > \pi_1'''$. Firm 1's optimal strategy is therefore to choose the point along Firm 2's reaction curve that is associated with the highest level of π_1 . This is given by point q_1^S , the point at which an isoprofit curve is tangent to Firm 2's reaction curve. This equilibrium is known as the **Stackelberg equilibrium**.²¹

Consider now the case when Firm 2 must pay an entry cost E before entering. Upon observing Firm 1's output level, Firm 2 will compare the entry cost to the prospective equilibrium profit in case it enters. The higher Firm 1's output, the lower Firm 2's expected payoff

²⁰Specifically, assuming that firms have the same constant marginal cost, Firm 2 enter unless Firm 1 chooses an output greater than the output under perfect competition.

²¹Stackelberg was the first to propose a variant of the Cournot model in which decisions are taken sequentially rather than simultaneously. See von Stackelberg, H, *Marktform und Gleichgewicht*, Vienna: Springer, 1934.




Figure 4: Stackelberg equilibrium.




Figure 5: Reaction curve with entry costs.

from entering (cf Exercise 15.6.b). Let q_1^L be the level of Firm 1's output such that Firm 2 is indifferent between entering and not entering, that is, q_1^L is such that $\pi_2(q_1^L, q_2^*(q_1^L)) = E$.

Assume that, when Firm 2 is indifferent between entering and not entering, it prefers not to enter. We can then plot Firm 2's new reaction curve to Firm 1's output choice, including not only its choice of output but its entry decision as well. Figure 5 depicts this new reaction curve. For values of q_1 lower than q_1^L , $q_2^*(q_1^L)$ is the same as before. The only difference from Firm 2's point of view is that it now must pay a positive entry cost E , but this has no influence on the post-entry output choice. If q_1 is greater or equal to q_1^L , however, then Firm 2 prefers not to enter the market, that is, $q_2^*(q_1^L)$ is zero.

As before, Firm 1 anticipates that Firm 2 will choose a point on its reaction curve. Consequently, Firm 1's optimal strategy is to choose the point in Firm 2's reaction curve that corresponds to the highest possible isoprofit curve. It can be seen from Figure 6 that this is given by q_1^L , precisely the output level beyond which Firm 2 finds it optimal not to enter. In words, Firm 1 finds it optimal to *increase the level of output with the sole purpose of deterring entry by Firm 2*.

Commitment, ex-ante optimality, and ex-post optimality²²

Until now, we have been a bit vague about the exact nature of the incumbent's initial move. Does Firm 1 set the level of capacity or does it set the output level instead? The distinction can

²²This section is adapted from Dixit, A., "The Role of Investment in Entry Deterrence," *Economic Journal* **90** (1980), 95–106.




Figure 6: Entry preemption.




Figure 7: Entry accommodation.

be very important. Suppose first that capacity constraints are not very important. Specifically, suppose that all of the production costs are incurred at the time of choosing output and that any level can be chosen at any time. In that case, even though Firm 1 is in the market before Firm 2, Firm 1 will be able to revise its output decision even after Firm 2 enters the market (if it does). For all practical purposes, what Firm 1 is doing in the equilibrium in Figure 6 is to announce its *intention* of producing output q_1^L in case Firm 2 enters. But then we may ask: is this a *credible* announcement by Firm 1?

In order for the announcement to be credible, it must be optimal for Firm 1 to implement what it announces when the time comes for making the relevant decision. Suppose that Firm 2, ignoring Firm 1's announcement, decides to enter and set output at the Cournot equilibrium level, that is, $q_2 = q_2^N$ (cf Figure 8). Then Firm 1's best strategy would be to ignore its pre-announcement and choose $q_1 = q_1^N$ instead. In fact, the same is true in the equilibrium with accommodation (Figure 4): instead of taking Firm 1's announced output as given, Firm 2 could simply enter and set output at the Cournot level.

Now consider an alternative interpretation of the model. Before Firm 2 decides whether or not to enter, Firm 1 must choose production capacity K_1 . Capacity costs are k per unit and *this cost is sunk*. That is, even if Firm 1 later chooses an output lower than installed capacity it will still not avoid the cost kK_1 . In addition to capacity costs, Firm 1 must also pay output costs cq_1 , that is, c per unit of output. Given that its capacity is K_1 , output will have to be less than capacity, that is, $q_1 \leq K_1$.

Let us consider the extreme case when capacity costs are very high and production costs are very low. This situation is depicted in Figure 9. The reaction curves $q_1^*(q_2)$ and $q_2^*(q_1)$ take into account both capacity and production costs.²³ If we were to consider production costs only, then, for Firm 1, we would obtain a reaction curve shifted rightwards, as in Figure 9. If Firm 1 has installed a large capacity, then this would be the relevant reaction curve. In fact, since capacity costs are sunk, only production costs (very low) matter, thus leading to a reaction curve far from the origin. This is all under the assumption that Firm 1 has a large capacity. Specifically, this is all under the assumption that Firm 1's optimal output would be lower than its capacity. The true ex-post reaction curve for Firm 1 is then the minimum of its capacity and the production-costs-only reaction curve, as depicted in Figure 9 and denoted by $q_1^*(q_2|K_1 = q_1^L)$: this means optimal output for Firm 1 given that Firm 2's output is q_2 and Firm 1 has previously installed capacity K_1 at the level q_1^L .

As can be seen from the Figure, even if Firm 2 were to enter and set a (reasonable) output level, it would still be optimal for Firm 1 to set output q_1^L , as announced. In other words, because production costs are so low with respect to capacity costs, it is optimal for Firm 1 to use all of its capacity. The preemption strategy is therefore a credible strategy.

Notice that, in addition to the assumption that capacity costs are high, we also need the assumption that they are sunk. If capacity costs were not sunk, then the first argument in this section would apply: Firm 2 could enter and set $q_2 = q_2^N$, and Firm 1's best reaction would be to sell capacity, recover the costs of investments previously made, and set $q_1 = q_1^N$ as well. We thus conclude that: Capacity preemption is a credible strategy only if capacity costs are high and sunk.




Figure 8: Non-credible preemption.

Contracts as a barrier to entry

The model below formalizes the discussion in the end of Section 15.1: the idea that a contract between an incumbent firm and a customer can create a barrier to entry.

Suppose that there is one buyer who is willing to pay up to one unit for the price of 1. The

²³Firm 2 will also have to pay capacity costs. In fact, suppose that the reaction curve $q_2^*(q_1)$ (in Figures 6 and 9) includes both production and capacity costs. This is the correct way of determining the equilibrium, for, when making its entry decision, Firm 2 should consider both capacity and output costs.




Figure 9: Credible entry preemption.

incumbent seller has cost $\frac{1}{2}$. The potential entrant's cost is unknown to buyer and incumbent. Both expect it to be a value uniformly distributed in the $[0, 1]$ interval. If entry takes place, the duopolists play a price-setting game, whereby the low cost firm sets a price just below the high-cost firm's cost.

Suppose there is no contract between incumbent and buyer. If the entrant's cost is greater than $\frac{1}{2}$, then no entry takes place and the incumbent sets the monopoly price 1. If the entrant's cost is less than $\frac{1}{2}$, then the entrant enters and sets price equal to the incumbent's cost, $\frac{1}{2}$. That is, if no contract is written, then the buyer's expected payoff is $\frac{1}{2}$ times zero (no entry) plus $\frac{1}{2}$ times $(1 - \frac{1}{2})$ (entry takes place), that is, $\frac{1}{4}$. The incumbent, in turn, receives an expected payoff of $\frac{1}{2}$ (no entry) times $(1 - \frac{1}{2})$, or $\frac{1}{4}$.

Consider now the following contract the incumbent offers the buyer. The buyer promises to purchase from the incumbent at a price $\frac{3}{4}$. If the buyer later wishes to switch to the entrant, then the buyer must pay the incumbent a penalty of $\frac{1}{2}$. We now show that the buyer is as well off, and the incumbent strictly better off, with this contract than without it. First notice that, despite the contract, entry still takes place when the entrant is very efficient. In fact, if the entrant's cost is less than $\frac{1}{4}$, then the entrant enters and sets price equal to $\frac{1}{4}$. Given this price, the buyer switches seller: it now pays $\frac{1}{4}$ (price) plus $\frac{1}{2}$ (penalty from breach of contract), a total of $\frac{3}{4}$ (the price required by the incumbent). Following this reasoning, we conclude that entry occurs with probability $\frac{1}{4}$, the probability that the entrant's cost is lower than $\frac{1}{4}$. The buyer's expected utility is $\frac{1}{4}$ times $(1 - \frac{1}{4} - \frac{1}{2})$ (entry takes place) plus $\frac{3}{4}$ times $(1 - \frac{3}{4})$ (entry does not take place), or simply $\frac{1}{4}$ (the same value as before). The incumbent, in turn, now gets $\frac{1}{4}$ times $\frac{1}{2}$ (entry takes place and the entrant receives the damage payment) plus $\frac{3}{4}$ times $\frac{3}{4} - \frac{1}{2}$ (entry does not take place), or $\frac{5}{16}$.

Not surprisingly, the loser in this process is the entrant. Without contracts, the entrant sells at a price $\frac{1}{2}$, whereas with the above contract the maximum the entrant can charge is $\frac{1}{4}$. Notice that, if the entrant's cost lies in the interval $[\frac{1}{4}, \frac{1}{2}]$, then entry does not take place even though the entrant is more efficient than the incumbent.

Mergers and firm value

The model below formalizes the discussion in Section 15.3.

The Cournot model with asymmetric firms illustrates how a merger may increase or decrease the value of non-merging firms. It can be shown that in a linear Cournot model where each firm i has marginal cost c_i , profits are given by

$$\pi_i = \left(\frac{a + \sum_{j \neq i} c_j - nc_i}{n+1} \right)^2 - F, \quad (15.1)$$

where $\sum_{j \neq i} c_j$ means summation of c_j for all j different from i , that is, the summation of marginal costs for all of firm i 's rivals.

Suppose we start from a situation with 3 firms, all with marginal cost $c_i = \bar{c}$. Each firm's profit is given by

$$\pi = \left(\frac{a - \bar{c}}{4} \right)^2,$$

where for simplicity we assume $F = 0$. Suppose that Firms 2 and 3 merge and that their new marginal cost is $\underline{c} < \bar{c}$. After the merger, Firm 1's profit is given by

$$\pi_1 = \left(\frac{a + \underline{c} - 2\bar{c}}{3} \right)^2.$$

(Notice that the number of firms is now 2.) Two situations are possible. First, if the merger between firms 2 and 3 does not imply large efficiency gains reflected in the level of marginal cost, then $\bar{c} \approx \underline{c}$ and $\pi_1 \approx (a - \bar{c})^2/3^2$, which is greater than the initial value (the denominator is smaller). If however \bar{c} is high (say, $\bar{c} \approx a/2$) and \underline{c} is much lower than \bar{c} (say, $\underline{c} \approx 0$), then $\pi_1 \approx 0$, which is lower than the initial value ($(a/8)^2 > 0$). In other words, depending on the effect of the merger on marginal cost, non-participating firms may gain or lose from the merger. If the merged firm increases its efficiency by decreasing marginal costs to a significant extent, then non-participating firms' value decreases as a result of the merger.

16 Research and development

The dynamics of R&D competition

The model below presents an alternative approach to the problem discussed in Section 16.2.

Suppose there are two firms: an incumbent monopolist, Firm 1, and a potential competitor, Firm 2. Both firms must decide how much to invest in R&D, r_i , $i = 1, 2$. The probability of success in R&D is given by $f(r_i)$. Success is independent across firms. Firms' profits are a function of the outcome of the R&D stage, as given in the main text. We assume that $\bar{\pi}^M > \underline{\pi}^M > \bar{\pi}^D > \pi^D > \underline{\pi}^D > 0$.

Each firm's expected payoff from the R&D investment is given by

$$\begin{aligned} \Pi_1 &= f(r_1) \left[f(r_2)\pi^D + (1 - f(r_2))\bar{\pi}^M \right] + \\ &\quad + (1 - f(r_1)) \left[f(r_2)\underline{\pi}^D + (1 - f(r_2))\underline{\pi}^M \right] - r_1 \\ \Pi_2 &= f(r_2) \left[f(r_1)\pi^D + (1 - f(r_1))\underline{\pi}^D \right] - r_2. \end{aligned}$$

The derivative with respect to r_i , $i = 1, 2$, is

$$\begin{aligned}\Pi'_1 &= f'(r_1) \left[f(r_2)\pi^D + (1 - f(r_2))\bar{\pi}^M \right] + \\ &\quad - f'(r_1) \left[f(r_2)\underline{\pi}^D + (1 - f(r_2))\underline{\pi}^M \right] - 1 \\ \Pi'_2 &= f'(r_2) \left[f(r_1)\pi^D + (1 - f(r_1))\underline{\pi}^D \right] - 1.\end{aligned}$$

It follows that the first-order conditions for profit maximization are

$$\begin{aligned}f'(r_1) \left[f(r_2)(\pi^D - \underline{\pi}^D) + (1 - f(r_2))(\bar{\pi}^M - \underline{\pi}^M) \right] &= 1 \\ f'(r_2) \left[f(r_1)\pi^D + (1 - f(r_1))\underline{\pi}^D \right] &= 1.\end{aligned}$$

The left-hand side of these equations gives the marginal benefit from R&D, whereas the right-hand side gives the marginal cost, one. Notice that, if R&D is not very drastic, then $\bar{\pi}^M \approx \underline{\pi}^M$ and $\bar{\pi}^D \approx \pi^D \approx \underline{\pi}^D > 0$. It follows that, for a given $r_1 = r_2 = r > 0$, the marginal benefit for Firm 1 is close to zero, whereas for Firm 2 it is positive. It can be shown that this leads to an equilibrium with $r_2 > r_1$. (Notice that, in the extreme when the above approximate equalities are exact, $r_1 = 0$ in equilibrium.)

17 Networks and standards

Diminishing marginal network effects

It is common to model network externalities as a utility function that includes network size as an argument. If utility is increasing in the size of the network, then we say that network externalities are positive. Normally, the increase in utility derived from a larger network is smaller the larger the network is; in other words, utility is normally a concave function of network size. What are the microeconomic foundations for such a utility function? In the next paragraphs, we present two simple models that illustrate how such a utility function can be arrived at. We consider both the cases of direct and indirect network externalities.

Let us start with the case of direct network externalities. Consider a total population of N email users and suppose that only a fraction λ actually have email. Each potential user would like to send m emails per period. Each of these emails can be sent to either one of two users. For example, suppose the user wants to ask a question and that there are two people who know the answer.

The utility from sending an email to the preferred destination is \bar{u} , whereas the utility from sending it to the second-best destination is \underline{u} , where $\bar{u} > \underline{u} > 0$ (for example, the first source will provide a better answer). If λ is the fraction of users who have email, it is also approximately the probability that a given user has email. The expected utility per email is therefore given by $\lambda\bar{u} + (1 - \lambda)\underline{u}$. The first term is the probability the preferred user has email times the respective utility. The second term is the probability that the preferred user does not have email but the second preferred does. Total per period average utility is finally given by

$$U = m \left(\lambda\bar{u} + (1 - \lambda)\underline{u} \right) = a(b - n)n,$$

where $n \equiv \lambda N$ is network size, and $a \equiv m \frac{u}{N^2}$, $b \equiv \frac{\bar{u}+u}{u} N$ are parameters. As can be seen, utility is an increasing, concave function of network size n . This is intuitive: when the number of email users is small, adding one more user to the network increases average user utility by more than when the new user is added to a large network (because the larger the network the greater the probability that the preferred user is already connected).

Let us now consider the case of indirect network externalities. Suppose that n consumers purchase a given operating system. Consumers do not derive any direct utility from using an operating system. Rather, utility comes from the quality of the applications software that is available. Specifically, the quality of applications software is a function of investment in software development: $b(r) = \beta r^\gamma$, where r is investment in software development and β, γ are parameters. Suppose that the gross benefit $b(r)$ is divided between consumers and software developer, so that the latter receives a per-consumer profit of $\alpha b(r)$ and the former a net benefit of $(1 - \alpha)b(r)$.

The optimal investment in software development maximizes $n \alpha b(r) - r$, where n is the total size of the installed base of operating software owners. The solution is given by $r^* = r^*(n)$. Plugging this back into the benefit function and simplifying, we conclude that each consumer receives a utility of

$$U = (1 - \alpha)\beta (\alpha \beta \gamma n)^{\frac{\gamma}{1-\gamma}},$$

which, under the assumption $\gamma < 1/2$, is an increasing, concave function of network size n .

Network effects and fulfilled-expectations equilibria

The model below formalizes the discussion in Section 17.1.²⁴

The demand for products subject to network externalities has some peculiarities which are not present in “normal” demand curves. The utility each consumer derives from the product depends on how many other consumers there are who purchase the same product — the size of the network of users. Or, to be more precise, demand depends on what each consumer *expects* the size of the network will be. The previous sentence points to an important element in the determination of demand under network effects: **consumer expectations**.

To illustrate this point, suppose there are one million consumers for a new technology subject to network effects. Each consumer’s valuation for the product (in dollars) is given by $v \cdot n$, where v is a parameter specific to each consumer and n is the size of the network (in million users). That is, the greater the value of n , the greater the valuation each potential buyer has for the product. The above valuation function implies that each consumer is willing to pay up to $v \cdot n^e$ for the product, where n^e is the *expected* size of the network.

Suppose that the value of v is uniformly distributed between zero and 1,000. Suppose moreover that consumers expect the size of the network to be 1m users, i.e., $n^e = 1$. Then the demand curve is given by the top curve in Figure 10. For example, if price is \$400, then a consumer with $v = 400$ is willing to pay exactly \$400, the product of $v = 400$ and $n^e = 1$. Consequently, all consumers with $v > 400$ are willing to pay more than $p = \$400$. Since v is uniformly distributed between zero and 1000, it follows that demand is given by .6m, the number of potential buyers with v greater than 400.

²⁴This section draws on “Network Effects,” a note to accompany the book by Shapiro, Carl, and Hal Varian, *Information Rules: A Strategic Guide to the Network Economy*, Cambridge, Mass.: Harvard Business School Press, 1998.

Suppose however that consumers expect network size to be $n^e = .5$ only. Then a consumer with $v = 400$ would only be willing to pay $v \cdot n^e = \$400 \times .5 = \200 . By contrast, a consumer with $v = 800$ would be willing to pay $\$800 \times .5 = \400 . If price is set, as before, at $p = \$400$, then all consumers with $v > 800$ would be willing to buy the product. Since v is uniformly distributed between 0 and 1000, demand is now given by 200,000, or .2m.

Similar calculations lead to a series of demand curves as in Figure 10. Each demand curve corresponds to a level of consumer expectations: the greater the expected network size, the greater demand is, for a given price. Specifically, if consumers expect a network of size n^e and a price p , then a consumer with v such that $p = v \cdot n^e$ is willing to pay as much as the product costs. We conclude that $v' = p/n^e$ is the value of v for such indifferent buyer. A fraction $1 - v'$ of the one million consumers is willing to pay more than price (since v is uniformly distributed). It follows that demand is given by $q = 1 - v' = 1 - p/n^e$. Substituting the values $n^e = 1, .5, .25$ for n^e we get the various demand curves in Figure 10.




Figure 10: Consumer expectations and demand curve.

In the short run, it is not unlikely that consumers' expectations are not precisely fulfilled. However, we would expect consumers to gradually adjust their expectations so that, eventually, expectations come close to the observed values. A **fulfilled-expectations equilibrium** is one where consumer expectations are exactly equal to the realized value of network size. In terms of model above, fulfilled expectations imply that $q = n^e$. Plugging in the equation above, we get $q = 1 - p/q$, which can also be written (if $q \neq 0$) as $p = q(1 - q)$.

Figure 11 depicts the fulfilled-expectations demand curve, that is, the demand curve corresponding to points where consumer expectations are exactly fulfilled. Notice that, for any price, $q = 0$ is always a point of the fulfilled-expectations demand curve. In fact, if consumers expect the network size to be zero, then each consumer is willing to pay $v \times 0 = 0$ for the product; that is, for any positive price demand is zero, which confirms the expectation that network size is zero. For prices greater than p'' (e.g., p') this turns out to be the only point in the demand curve. However, for prices below p'' , in addition to the zero-demand point, we find two other points of fulfilled expectations. For example, $p = p'''$ implies that both $q = n'$ and $q = n'''$ belong to the demand curve. We conclude that network effects may imply multiple demand levels for a given price. Which value takes place depends on consumers' expectations regarding network size.

What can we say about this “chicken-and-egg” problem? For example, if $p = p'''$, which of the three network sizes $(0, n', n''')$ will result? One thing we can say is that the value n' is unlikely to result. The reason is the following. Suppose that the value of q is slightly greater than n' . Then the valuation for the product, at the margin, is greater than price (the demand curve is higher than p'''). We would expect this to lead additional consumers to join the network, implying an even greater divergence between the actual network size and the initially expected network size (n'). Likewise, if the value of q is slightly lower than n' , consumers are willing to pay less than price. We would then expect consumers to leave the network, implying an even greater divergence between the actual network size and the initially expected network size (n'). In summary, we would expect the value n' not to be *stable*. More generally, the dashed section of the demand curve (the increasing portion) corresponds to unstable points, points which we would not expect to find in equilibrium.

What about the choice between $q = 0$ and $q = n'''$? Both are stable points. Specifically, suppose that q is slightly greater than zero but very small. As the demand curve in Figure 11 shows, it would still be the case that, at the margin, the value of the product is less than price: the few consumers who bought the product would then leave the network and network size would go back to zero, one of the fulfilled-expectations levels. Likewise, suppose that q is slightly less than n''' . As the demand curve in Figure 11 shows, at the margin consumers would be willing to pay more than price: some consumers would then join the network, thus reestablishing the equality between q and the expected network size n''' . In summary, both $q = 0$ and $q = n'''$ are stable points and there does not seem to be much one can say about which one is more realistic.

Now suppose that p is very low. This means that even a short deviation in the value of q would move us to the right of the dashed line. That is, even a small perturbation in the value of q would lead to a point where the demand curve is higher than price. As seen above, this would imply a tendency for demand to increase even further — to the large-network equilibrium, in fact.

The above analysis suggests two points. First, convergence to the high-equilibrium depends on passing the threshold given by the dotted line. Once that threshold is crossed, demand will continue increasing in a self-reinforcing process that ends in the large-network equilibrium. This threshold level is usually referred to as the **critical mass** of buyers that leads to the buildup of the network. The second point is that the lower price is the greater the likelihood that the threshold is crossed, i.e., critical mass is achieved.

These two points have a number of implications. In a competitive market, where price depends primarily on cost considerations, and technical progress drives costs down over time, we would expect the initial equilibrium to be a high price and a very small or non-existing network (such as in the point $(p = p''', n = 0)$ in Figure 11). As time passes, cost goes down and so does price. Eventually, price crosses the threshold p'' , below which there are two stable equilibria. At some stage (e.g., when $p = p'''$), due to random perturbations in demand, network size crosses above the critical level given by the dotted line in Figure 11 and demand converges to the high-network equilibrium n''' . From then on, additional downward movements in price result in slight increases in network size.




Figure 11: Network externalities and fulfilled-expectations equilibrium.