Journal of Industrial Organization Education

Volume 2, Issue 1

2008

Article 3

Price Discrimination

Sylvain Weber, *University of Geneva, Department of
Economics*

Cyril Pasche, *University of Geneva, Department of
Economics*

Recommended Citation:

Weber, Sylvain and Pasche, Cyril (2008) "Price Discrimination," *Journal of Industrial Organization Education*: Vol. 2: Iss. 1, Article 3.

DOI: 10.2202/1935-5041.1020

Price Discrimination

Sylvain Weber and Cyril Pasche

Abstract

This lecture, PowerPoint slide show, and Stata ado-file offer an introduction to price discrimination, with an emphasis on third-degree price discrimination. Graphical intuition and mathematical analysis make this lecture appropriate both for undergraduate and graduate students. After explaining the functioning of a simple monopoly, and discussing first and second-degree price discrimination, the lecture deals with third-degree price discrimination. Everyday examples, graphics, and a mathematical treatment are offered. The lecture and the accompanying PowerPoint slide show should take about 90 minutes to go through.

KEYWORDS: monopoly, price discrimination, kinked demand

Author Notes: University of Geneva, Department of Economics, 40 Bd du Pont-d'Arve, 1211 Geneva 4, Switzerland. A Stata ado-file solving the third-degree price discriminating monopoly problem is available for download on the website of the Journal of Industrial Organization Education. We thank Giovanni Ferro Luzzi, Nicolas Maystre and Marcelo Olarreaga for highly valuable suggestions. We also thank Jeffrey M. Perloff for very helpful comments that led to revisions and clarifications.

Erratum

On April 25 2008, the citation for this article was changed from

Weber, Sylvain and Pasche, Cyril (2007) "Price Discrimination," *Journal of Industrial Organization Education*: Vol. 2 : Iss. 1, Article 3.

to

Weber, Sylvain and Pasche, Cyril (2008) "Price Discrimination," *Journal of Industrial Organization Education*: Vol. 2 : Iss. 1, Article 3.

Introduction

This lecture deals with price discrimination in general and third-degree price discrimination in particular. It is split into seven sections and covers the essentials of price discrimination. Section 1 offers examples of price discrimination. Section 2 recalls the conditions that must be met for a simple monopoly to form and be sustained over the long term. It also offers a brief explanation of the functioning of a monopoly. Section 3 lists the additional conditions that are necessary for a monopoly to be able to price discriminate. Section 4 briefly presents, both graphically and mathematically, first-degree price discrimination. Section 5 focuses on second-degree price discrimination. Section 6 presents the third-degree price discrimination model and offers an in-depth mathematical development. Contrary to most textbooks, we offer a presentation that is not based on a constant marginal cost and is thus more general than other texts. Graphical intuition and mathematical analysis render the discussion accessible for both undergraduate and graduate students. Section 7 summarizes the main points of the lecture.

A Stata ado-file solving the third-degree price discriminating monopoly problem numerically and graphically is provided. The program's functional and installation instructions are explained in the Appendix.

1 Price Discrimination: A Widespread Practice

Price discrimination happens when units of the same product are sold at different prices, either to the same consumer or to different consumers. According to the classification of Pigou (1920), price discrimination may take three forms. First-degree price discrimination is perfect price discrimination. The producer is able to charge whatever the market will bear and thus captures the entire consumer surplus. Second-degree price discrimination involves selling larger quantities of the same product at a lower unit price. Third-degree price discrimination takes place when different prices are charged for the same product to different consumers.

The monopoly's aim, when practicing price discrimination, is to extract most of the consumer surplus. Price discrimination always increases the profit of the monopoly, and, in the vast majority of cases, decreases the welfare of consumers and society as a whole. United States federal law legislates price discrimination through the Robinson-Patman Act of 1936 (also called the Anti-Discrimination Act). The full Act consists of a single, long, confusing sentence with 370 words and 30 commas. To get an idea of its complexity,

let us enumerate the conditions that must be met to bring the Act into play. There must be (1) two or more consummated sales, which (2) must occur close in time, and (3) involve commodities (not services). The goods must be (4) of like grade and quality, (5) be sold at different prices by the (6) same seller, to (7) two or more different purchasers for (8) use, consumption, or resale within the United States, which may (9) result in competitive injury. The Robinson-Patman Act has a fairly narrow scope.¹

Price discrimination occurs very frequently. Nearly all firms with market power attempt to engage in some form of price discrimination and even governments practice it. The following real-life examples illustrate these facts. A U.S. passport for people under 16 costs USD 85; for people aged 16 and above, it costs USD 100. With a three night stay at “The New York Palace” you get a complimentary fourth night, i.e., people who stay longer get charged less per night. American and Canadian students pay USD 165 for the sixth edition of Greene’s *Econometric Analysis*; the international version costs only the equivalent of USD 100.² Institutions pay USD 247 for a one year subscription to *The Journal of Industrial Economics*, individuals pay USD 78 and students USD 30. Early-bird discounts, and peak and off-peak pricing are other examples of price discrimination.

2 Simple Monopoly

A simple monopoly occurs when: (1) a single firm is the sole provider of a product, (2) barriers to entry ensure that no concurrent will enter the market, and (3) there are no close substitutes to the monopolized product.

In the case of a simple monopoly, every consumer pays the same price for any amount of a given product. Consumers with a willingness-to-pay above the market price purchase the product and obtain a surplus; those with a willingness-to-pay below the market price don’t purchase the product. Prescription drugs sold within a country represent a good example of a simple monopoly.

To maximize profit, the monopoly produces a quantity such that the marginal cost of production is equal to the marginal revenue, as in a competitive market. The difference between a monopoly and a competitive firm comes

¹See Stancil (2004) for a comprehensive guide of the Robinson-Patman Act.

²International editions are not sold within the United States or Canada, and the colophon of such books even precise that: “If you purchased this book within the United States or Canada you should be aware that it has been wrongfully imported without the approval of the Publisher or the Author.”

from the shape of the marginal revenue. The marginal revenue is constant and equal to the market price for a competitive firm, whereas it is below the market price and decreasing for a monopoly. This difference stems from the fact that a monopoly can choose a point along the demand curve and is thus considered a price-maker while a competitive firm is a price-taker. Because of the law of demand, a simple monopoly willing to sell an additional unit of product has to lower the price, so that marginal revenue is lower than market price.

Mathematically, maximizing the profit of a monopoly implies equalizing the derivative of the profit function to zero:

$$\pi(Q) = TR(Q) - TC(Q) = P(Q) \cdot Q - TC(Q)$$

$$\begin{aligned} \frac{\partial \pi(Q)}{\partial Q} &= 0 \\ \frac{\partial P(Q) \cdot Q}{\partial Q} - \frac{\partial TC(Q)}{\partial Q} &= 0 \\ P(Q) + \frac{\partial P(Q)}{\partial Q} Q - MC(Q) &= 0 \\ P(Q) \cdot \left(1 - \frac{1}{|\varepsilon_p^d|} \right) &= MC(Q) \end{aligned} \tag{1}$$

When the optimal quantity (Q_{SM}^*) has been determined, the price is chosen by projecting it in the demand curve. The monopoly thus charges the maximal possible price given the market demand. Using 1 to write the equilibrium price, we get:

$$P_{SM}^* = \frac{MC(Q_{SM}^*)}{1 - \frac{1}{|\varepsilon_p^d|}} \tag{2}$$

where $MC(Q_{SM}^*)$ is the value of the marginal cost for the optimal quantity and ε_p^d is the price elasticity of demand at equilibrium. The monopoly always operates in the elastic part of the demand curve ($|\varepsilon_p^d| \geq 1$).

It is clearly seen in the price equation (2) that the monopoly power is constrained by the market demand. If the demand is perfectly elastic ($|\varepsilon_p^d| = \infty$), then the monopoly price is equal to the marginal cost of production, which also corresponds to the perfectly competitive price. In such an extreme case, the monopoly power is absolutely nil. When the price elasticity of demand diminishes (in absolute value), the monopoly power increases and price moves

away from the marginal cost. In other words, the steeper the demand curve, the more a monopoly is able to raise the price.

On slide 5 we represent the demand curve (D) faced by a simple monopoly, its marginal revenue (MR) and its marginal cost (MC). The monopolist produces Q_{SM}^* and sells at price P_{SM}^* . In a competitive market, output would be Q_{CM} and price P_{CM} . The deadweight loss (lowering of the total surplus) due to the monopoly is given by the shaded triangular area.

3 Conditions for Price Discrimination

Most monopolies face consumers with different willingness-to-pay. To further increase their profits, they can consequently charge different prices to different consumers.

For a monopoly to be able to price-discriminate, the following conditions must arise, in addition to the three that were stated in section 2: (1) there exist several types of consumers, who have different sensitivities to price; (2) arbitrage is not possible, in the sense that consumers who are charged a lower price cannot resell the product to consumers facing a higher price, and (3) each consumer can be identified. Condition (3) is not required for second-degree price discrimination. When these elements are combined, a monopoly will be able to take advantage of its position by charging different prices to different groups of consumers.

For example, the Alcatraz cruise is charged differently for seniors (USD 23.25) and adults (USD 24.50), because: (1) adults are more willing/able to pay than seniors, (2) distinguishing adults from seniors is simple, and (3) arbitrage is impossible. In contrast, when buying an apple at the local grocery store, consumers all pay the same price, simply because arbitrage is probable.

4 First-Degree Price Discrimination

First-degree price discrimination is also known as perfect price discrimination. Each consumer is charged the maximum amount they are willing to pay for each unit. The consumer surplus is nil, because it is totally reaped by the producer. Selling an extra unit does not require lowering the price of the previously sold units, since each one may be sold at a different price. Demand and marginal revenue curves are therefore confounded, and the profit maximizing condition (marginal revenue equals marginal cost) yields the same output as a competitive market. If a monopoly is able to practice perfect discrimination,

the outcome will thus be Pareto efficient, which is a quite counter-intuitive result. Transaction and research costs involved in finding the willingness-to-pay of each consumer make this case extremely unlikely in the real world.

We illustrate first-degree price discrimination by examining how Dr. Bernard fixes his rates. Dr. Bernard is the only medical doctor in his Montana hometown of Kalispell. He has lived around Kalispell all his life and personally knows each of his patients. He thus knows how much each one is willing to pay for his services. Dr. Bernard is a smart man and knows he can make more money by charging each patient a different fee corresponding to what each one is susceptible to pay, instead of applying a fixed rate. The higher the willingness-to-pay of the patient, the higher the fee.

In the figure of slide 9, outcome Q_F^* is produced by the first-degree price discriminating monopoly. The monopolist makes a profit on every unit produced equal to the difference between price and marginal cost. We can see that consumer surplus disappears completely, as the producer captures the whole surplus. There is, however, no deadweight loss in this case. Interestingly, the first-degree discriminating monopoly produces just as much as if it were in a competitive market, since $Q_F^* = Q_{CM}$.

The profit of a first-degree price discrimination monopoly is given by the sum of the different revenues of each unit minus the costs for producing these units:

$$\pi(Q_F^*) = \int_0^{Q_F^*} [P(Q) - MC(Q)]dQ - FC = \int_0^{Q_F^*} [P(Q)]dQ - TC(Q_F^*)$$

where FC are the fixed costs. In slide 9, profit is given by the shaded area minus the fixed costs.³

5 Second-Degree Price Discrimination

Second-degree price discrimination is also known as nonlinear pricing: prices vary across units, but not across people. The usual example of second-degree price discrimination is quantity discounts, when larger quantities of the same product are sold at a lower unit price. Producers use this technique when they know different consumer groups exist, but have no way of observing them. The complete extraction of consumer surplus is not possible as it was in the first-degree price discrimination, but the producer can still design some price

³The difference between consumer surplus and profit is equal to the fixed costs: $\pi = CS - FC$.

scheme to force consumers to reveal their type. Second-degree price discrimination is a more realistic case of price discrimination. It is often practiced by public utilities, as the price per unit of water and electricity often depends on how much is consumed. Consumers self-select themselves into consumption categories.

We illustrate second-degree price discrimination by examining how much beer John and Richard order at the Moab Brewery for their Saturday night party. The Moab Brewery is the sole microbrewery in south-east Utah. Just by seeing John and Richard, the brewer has no effective way of telling them apart and knowing who has a high or a low demand for beer. To get around this problem, the brewer offers a pint of oatmeal stout for USD 2.50 and a pitcher (half a gallon, which is equivalent to four pints) of the same beer for USD 7.50. The unit price of a pint is thus USD 1.875 when buying a pitcher. The pint is directed to low demand consumers, and the pitcher to high demand consumers.

In the example illustrated in slide 11, John orders a pint and Richard a pitcher. John's total willingness-to-pay for one pint is given by zone A. It is at least equal to USD 2.50 since John has decided to buy a pint. For simplicity, we assume that A is equal to USD 2.50 so that John's consumer surplus is zero. John will not want to buy a pitcher, since his willingness-to-pay for a pitcher (A+B) is lower than USD 7.50. Let us now consider Richard's choice. He could buy one pint for USD 2.50 and make a surplus of C. Since his demand is high, he is also susceptible to pay more than USD 7.50 for a pitcher. What is the criterion he will consider for choosing between the pint and the pitcher? Richard will obviously opt for the solution that gives him the greatest surplus. The maximal price he is willing to pay for a pitcher is thus given by area A+D, which leaves him with a surplus exactly equivalent to what he has with a pint (B). If $A+D = \text{USD } 7.50$ (or just a little more so that he is not indifferent between the two options), Richard will buy a pitcher. With such a price scheme, consumers self-select into their true category.

6 Third-Degree Price Discrimination

Third-degree price discrimination consists of distinguishing several groups of consumers with different demands in order to charge them different prices. Producers discriminate between groups of people, but not across units, as consumers are charged different prices but each one faces a constant price for all units of output purchased. It is used to take advantage of the fact that different groups of people have different demand functions. The usual

example of third-degree price discrimination is student discounts at the movies or at sports events. Third-degree price discrimination is also practiced by pharmaceutical firms who sell the same drugs in different countries at different prices. For example, Tamiflu⁴ is sold for USD 65 in the United States, AUD 45 in Australia, and CHF 85 in Switzerland.⁵

Third-degree price discrimination is a complex issue and represents the core of today's lecture. We begin by presenting the non-discriminating model with a market segmented between several consumer groups and then switch to the discriminating one. In a quest of simplicity, we will assume the world is made of two markets: Northern countries with a high ability to pay and Southern countries with a low ability to pay.

6.1 Simple Monopoly on a Segmented Market

Three possibilities may explain why a monopolist facing different consumer types does not charge different prices: (1) the producer simply ignores the fact that several consumer groups with different attitudes are buying its product, (2) it is not possible to identify the consumers, or (3) consumers can resell the product. In these cases, the monopolist will act in the “standard way”; producing the quantity equalizing its marginal revenue with its marginal cost and selling this quantity at the maximal possible price given the market demand.

The construction of the aggregate demand as well as the corresponding marginal revenue is shown on slide 13. The aggregate demand considered by the simple (non-discriminating) monopoly is given by the (horizontal) sum of the different group demands. With two consumer groups, the aggregate demand comprises an “inwards” kink. Since the slope of the demand becomes suddenly “flatter” at Q_{kink} , the marginal revenue curve is broken and jumps up.

The fact that the marginal revenue jumps up implies three different cases: the marginal cost may cross the marginal revenue (1) only once and in its upper segment, (2) twice or (3) only once and in its lower segment.

In case (1), the equilibrium price will be too high for the group with a low willingness-to-pay, so that only individuals with a high willingness-to-pay consume. In case (3), both types of consumers will be served by a single-price monopoly.

⁴Tamiflu is the most effective prescription drug against the H5N1 bird flu, which is currently the worst pandemic threat.

⁵On April 3, 2008, AUD 1 was worth USD 1.10, and CHF 1 was worth USD 1.01.

The most interesting case (2) occurs when there are two intersections between the marginal revenue and the marginal cost, and therefore two extremums for the profit function. This case is illustrated in slide 14. The first of these extremums (Q_A) implies a price too high for the group of consumers with a low willingness-to-pay, and consequently only consumers from the high willingness-to-pay group buy the product. The second extremum (Q_B) is reached with a lower price allowing both groups to consume and a higher quantity traded. The monopolist has to compare the two extremums (Q_A and Q_B) to decide which one leads to the highest profit.

Graphically, the reasoning is as follows. Consider moving from point A to point B. Between point A and the discontinuity, the marginal profit is negative since the marginal cost outweighs the marginal revenue. A loss is therefore suffered on every unit produced between Q_A and Q_{kink} . At the right of the discontinuity, marginal revenue is higher than marginal cost and marginal profit is thus positive between Q_{kink} and Q_B . We need to compare the two triangular areas defined by the marginal cost and the marginal revenue curve to find where the monopoly profit is globally maximized. The first triangle represents a loss and the second a gain, so that the units between Q_A and Q_B will only be produced if the second triangle is larger (or equal) than the first.

A direct consequence of the discontinuity of the marginal revenue is that a slight modification of the marginal cost may cause a great change in the quantity produced. Indeed, the monopolist having a marginal cost equalizing the two triangular shaded areas of slide 14 is indifferent between producing Q_A or Q_B . An infinitesimal decrease of the marginal cost would lead to a quantity greater than Q_B whereas a tiny increase would lead to an equilibrium quantity smaller than Q_A .

6.2 Third-Degree Price Discriminating Monopoly

Let us now consider that the monopoly is able to practice third-degree price discrimination, i.e., that all the conditions enumerated in section 3 are met. As usual, the producer's objective is to maximize profit, so that marginal revenue and marginal cost are equalized. What changes compared to the simple monopoly is the marginal revenue. When the monopoly does not price-discriminate, we look for the market demand first and retrieve the marginal revenue afterwards. For the discriminating monopoly however, one has to sum the marginal revenues of the different consumer groups horizontally to obtain the aggregate marginal revenue. Retrieving an aggregate demand for the discriminating monopoly is in fact not possible since total output is sold at several different prices.

The mechanism of the third-degree price discriminating monopoly can be explained in three simple steps:

1. The firm chooses its total output by equating its aggregate marginal revenue to its marginal cost:

$$MR_{tot}(Q^*) = MC(Q^*) = MR_{tot}(Q^*) \quad (3)$$

2. The allocation of the total output between the different consumer groups is made by equalizing the marginal revenue of each market segment to the *value* of the marginal cost for the optimal output Q^* , which is also equal to the value of the total marginal revenue because of (3):

$$MR_i(Q_i^*) = MC(Q^*), \quad i = 1, 2, \dots, N$$

(4)

where $\sum_{i=1}^N Q_i^* = Q^*$

3. The optimal prices charged to the different consumer groups are found by replacing the quantities in the demand functions:

$$P_i^* = P_i(Q_i^*) \quad (5)$$

Let us deal with the mathematical resolution of the problem using linear demand functions. The inverse demand of segment i is $P_i(Q_i) = a_i - b_i Q_i$, and we thus have $MR_i(Q_i) = a_i - 2b_i Q_i$. Without loss of generality, consider the market is composed by two groups of consumers (e.g., North and South) and let $a_N > a_S$. Since marginal revenues have to be equalized across market segments to maximize profit (equation 4), one has to sum them horizontally in order to obtain the overall marginal revenue of the discriminating monopoly. To do so, we first need to express Q_i as a function of MR_i :

$$Q_i(MR_i) = \frac{a_i}{2b_i} - \frac{1}{2b_i} \cdot MR_i, \quad i = N, S.$$

We are then able to sum the quantities, and using the fact that equation (4) must hold in equilibrium we get:⁶

$$Q_{tot} = \begin{cases} Q_N = \frac{a_N}{2b_N} - \frac{1}{2b_N} \cdot MR_{tot} & \text{if } MR_{tot} > a_S \\ Q_N + Q_S = \frac{a_N b_S + a_S b_N}{2b_N b_S} - \left(\frac{b_N + b_S}{2b_N b_S}\right) \cdot MR_{tot} & \text{if } MR_{tot} \leq a_S \end{cases}$$

⁶In what follows, $MR_{tot}(Q^*)$ is noted MR_{tot} for the sake of clarity.

Reversing once again and simplifying, we finally obtain the function of the aggregate marginal revenue for the discriminating monopoly:

$$MR_{tot} = \begin{cases} a_N - 2b_N \cdot Q_{tot} & \text{if } Q_{tot} < \frac{Q_{kink}}{2} \\ \frac{a_N b_S + a_S b_N}{b_N + b_S} - 2 \frac{b_N b_S}{b_N + b_S} \cdot Q_{tot} & \text{if } Q_{tot} \geq \frac{Q_{kink}}{2} \end{cases} \quad (6)$$

where $Q_{kink} = \frac{a_N - a_S}{b_N} \cdot \frac{Q_{kink}}{2}$ is the quantity at which the aggregate marginal revenue is kinked. One can check that there is no discontinuity in the marginal revenue by replacing $Q_{tot} = \frac{Q_{kink}}{2}$ in the two portions of the function, which gives a_S in each case. There is only a change in the slope of the marginal revenue at this point, but no discontinuity. Any marginal cost function therefore intersects the marginal revenue once and once only, which means that after having written the aggregate marginal revenue, there is no ambiguity in choosing the total quantity to be produced.

One needs to be careful, however, concerning the segment of the marginal revenue to be equalized with the marginal cost function. Concretely, one can begin to compute the aggregate quantity using the second portion of equation (6), i.e., assuming that both types of consumers will be served in equilibrium. If the optimal output level is found to be smaller than $\frac{Q_{kink}}{2}$, it means that only northern consumers will be served. The computation of the optimal then has to be re-done using the first portion of equation (6), which is simply the marginal revenue of the first market segment. In this latter case, the monopolist behaves like a simple monopoly only active on one market segment, since it is not profitable to sell on the other one.

Once the optimal total output is found, the monopolist still has to choose how to distribute it between the market segments. To maximize profit, the monopolist equalizes the marginal revenues in each market segments. As long as the marginal revenues differ between market segments, the monopolist can increase profit by reallocating units of output from the segment where marginal revenue is low to the segment where it is high. Since total output remains unchanged during this process, total costs do not change, but total revenue increases thanks to the reallocation. The profit thus rises. Practically, one has to equate the *value* of the marginal cost for the optimal total output with the marginal revenue function of each segment in order to find the optimal distribution of quantities. We stress that the calculation has to be made using the *value* of the marginal cost and not with its *function*. This latter procedure would imply that each market segment is treated individually and that costs depend on the quantity offered on one segment, whereas market segments are actually linked by the fact that costs depend on the sum of all quantities.

The final step is to find the optimal prices, which is done simply by re-

placing the optimal quantities in the inverse demand of each market segment. The monopolist's profit can eventually be computed:

$$\begin{aligned}\pi(Q_N^*, Q_S^*) &= TR_N(Q_N^*) + TR_S(Q_S^*) - TC(Q^*) \\ &= P_N^* \cdot Q_N^* + P_S^* \cdot Q_S^* - TC(Q^*)\end{aligned}\tag{7}$$

where $Q^* = Q_N^* + Q_S^*$. This final equation emphasizes the fact that both segments are related. The total cost has to be computed at the aggregate market level, so that individual segment profit cannot be obtained. What can be (and has to be) computed separately are the total revenues on each market segments, since prices are not the same.

A graphical representation of the third-degree price discriminating monopoly is provided in slide 17. Panels A and B represent the demands of two consumer groups. Knowing the demand curves of the two segments, one can derive the marginal revenues and construct the aggregate marginal revenue on panel C, which represents the entire market. The optimal total output (Q^*) is found at the intersection of the marginal cost and the aggregate marginal revenue. Projecting the value of the marginal cost for the optimal quantity ($MC(Q^*)$) on panels A and B (thin dashed line) shows the optimal distribution of the total output between the consumer groups (Q_N^* and Q_S^*), and finally allows for the determination of segment prices (P_N^* and P_S^*).

6.3 From the Simple to the Discriminating Monopoly

In order to complete this lecture, we compare what happens if a non-discriminating monopoly suddenly becomes able to practice price discrimination. Our discussion relies on slide 19.

Before going further, let us warn the reader about a possible confusion. Kwon (2006) states that plotting the aggregate demand “can confuse students into thinking that the simple monopoly's demand curve (D_{SM}) is the demand curve corresponding to the price-discriminating monopoly's marginal revenue curve, which is not because the total revenue of the price-discriminating monopoly cannot be derived from D_{SM} ”. We nevertheless believe it is useful since it allows one to visually compare how quantities and prices change when moving from a simple (non-discriminating) monopoly to a discriminating one (or inversely). In order to make it as clear as possible that the aggregate demand on panel C is in fact not related to the discriminating monopoly, we draw it as a dashed line and label it D_{SM} for “Demand of a Simple Monopoly”. Kwon (2006) derives the aggregate average revenue curve of the discriminating monopoly corresponding to its aggregate marginal revenue. The curve lies

above the aggregate demand of the simple monopoly. We do not enter into the details of the construction of this curve since this is beyond the scope of this lecture and invite the interested reader to consult this reference.

When moving to a discriminating monopoly, the broken marginal revenue of the simple monopoly suddenly becomes continuous and kinked at point I . In the example of slide 19, the total quantity produced by a non-discriminating monopoly would be the same as in the discriminating case. The price maximizing the profit of a simple monopoly would be P_{SM}^* . Projecting this price on panels A and B shows that it is somewhere between P_N^* and P_S^* , the prices charged by a discriminating monopoly. The price charged to the group with the more inelastic demand (D_N) is higher with price discrimination than without, and the quantity sold to this group is consequently lower under price discrimination. This result is perfectly intuitive: consumers with a more inelastic demand are less price-sensitive and it is optimal for the monopolist to charge them more. This result can be demonstrated rigorously by expressing the profit-maximizing condition (4) in the following form:

$$MC(Q^*) = P_N^* \cdot \left(1 - \frac{1}{|\varepsilon_{p_N}^d|}\right) = P_S^* \cdot \left(1 - \frac{1}{|\varepsilon_{p_S}^d|}\right) \quad (8)$$

where $\varepsilon_{p_i}^d$ is the value of the price elasticity of demand at equilibrium for market segment $i = N, S$. Equation (8) shows there is an unambiguous relationship between the price and the elasticity of demand for each market segment. If the demand in the South is more elastic than in the North ($|\varepsilon_{p_S}^d| > |\varepsilon_{p_N}^d|$), then price will be lower in the South ($P_S^* < P_N^*$).

Concerning the level of total output, two interesting cases have to be considered when the monopoly moves from a single price for the whole market to third-degree price discrimination.⁷ We use slide 14 to illustrate what happens. Firstly, if the non-discriminating monopoly maximizes its profit at an output level located on the second portion of its marginal revenue (like Q_B), there will be no change in the total output. Since both consumer groups are already served under uniform pricing, the effects of price discrimination would be as follows: price would increase for consumers with the more inelastic demand and decrease for those with a more elastic one, with quantities moving in the opposite direction. Secondly, if the quantity maximizing the monopolist's profit is located on the first portion of the marginal revenue (like Q_A) without discrimination, it means that the consumer group with low willingness-to-pay is not served. Price discrimination allows those consumers to be served, so

⁷See Battalio & Ekelund (1972) for an exhaustive analysis of output changes under third-degree price discrimination.

that optimal total output will suddenly jump to a larger level (Q_B), without any change in consumer demands nor in production costs. Consumers in the South are thus always better off with price discrimination. For consumers in the North, the move to price discrimination implies changes that depend on the slope of the marginal cost:

- If the marginal cost is upward sloping, their price increases and the quantity they consume decreases.
- If the marginal cost is constant, nothing changes.
- If the marginal cost is downward sloping, their price goes down and the quantity they consume increase.

In all the cases, the monopoly's profit will necessarily increase when moving from a single price to third-degree price discrimination. Though not always feasible, price discrimination is always profitable for the producer, since at worst, the monopolist can always charge a uniform price. Formally, the argument can be demonstrated as follows. For a monopolist charging a single price, the marginal revenues differ between consumer groups. A simple reallocation of the total output from a segment where marginal revenue is low (North) to another segment where it is high (South) will improve the producer's position. Total costs remain unchanged since total output is constant, but revenues increase thanks to the reallocation. As a result, profit rises.

7 Summary

In this lecture, we have explored an important and widespread economic practice: price discrimination. For a firm to price discriminate successfully, three conditions must be met: the firm must have market power; consumers must differ in their price sensitivity; and resales must be prevented. It is always profitable for a firm to practice price discrimination, since it allows the producer to reap some of the consumer surplus.⁸

Price discrimination may take three different forms. First-degree price discrimination is perfect price discrimination. Each consumer is charged the maximum price they are willing to pay, and the producer captures the entire consumer surplus. Second-degree price discrimination involves selling larger

⁸Our study of the welfare effects of price discrimination is quite elementary in this lecture. A more exhaustive analysis is provided by Varian (1989).

quantities of a same product at a lower unit price. Third-degree price discrimination consists in distinguishing several groups of consumers with different demands in order to charge them different prices.

A third-degree price discriminating monopoly allocates the total output between groups of consumers so that the marginal revenues are equalized across market segments. Consumers are thus charged a price that is related to the price elasticity of their demand: consumers with a more inelastic demand are charged more.

The total output produced by a third-degree price discriminating monopoly is always equal or greater than what is produced by a simple monopoly, even if both are serving the same consumers. Third-degree price discrimination leads to a greater output when consumers that did not initially consume with a single price for the whole market are now served. An interesting case occurs when the marginal cost of the producer is decreasing and one market segment is not served by a single-price monopolist. It is then possible to show that moving to third-degree price discrimination will have a positive effect for every market participant.

Appendix: Stata Command

As a complement to this lecture, we provide a command that solves the problem of a third-degree price-discriminating monopoly. This should prove especially useful to lecturers scratching their head to concoct new exercises in order to end up with an original exam each year.

The command is implemented as a Stata ado-file.⁹ We decided to use this software instead of a more widespread one like Microsoft Excel because the capabilities of Stata are far more developed, especially concerning graphical representations. The ado-file runs under release 8.0 or later of Stata.

The ado-file is called *tdpd* for “Third-Degree Price Discrimination”. It solves the problem of a monopolist facing two groups of consumers, both in the discriminating and the non-discriminating cases. It therefore computes the optimal total output to produce, the allocation of this output between the groups of consumers, prices charged to each group, profit of the monopolist, *asf*. Graphical representations are provided for the two types of monopolies. A help file is provided with the command.

⁹Both the ado-file and its companion help file are available for download on the website of the *Journal of Industrial Organization Education*. To make it work, one should save the files in a directory called: “C:\ado\plus”. For more information about ado-files, see Stata User’s Guide.

The required inputs are the parameters of the two linear demands and the parameters of the monopoly's total cost function. Denoting the demand functions of the two market segments $Q_i(P_i) = a_i - b_i P_i$, $i = 1, 2$ and the monopoly's total cost function $TC(Q) = c_1 + c_2 Q + c_3 Q^2$, the syntax of the command is:

```
tdpd    a1 b1 a2 b2 c1 c2 c3
```

For example, if the demand of group 1 is $Q_1(P_1) = 100 - P_1$, the demand of group 2 is $Q_2(P_2) = 60 - P_2$, and the monopoly's total cost function is $TC(Q) = 10 + 5Q + 0.25Q^2$, one only has to type the following line in the Stata browser:

```
tdpd    100 1 60 1 10 5 .25
```

The software will do the rest.

References

- Battalio, R. C. & Ekelund, R. B. (1972), 'Output Change under Third Degree Discrimination', *Southern Economic Journal* **39**(2), 285–290.
- Kwon, Y. (2006), 'Third-Degree Price Discrimination Revisited', *Journal of Economic Education* **37**(1), 83–92.
- Pigou, A. C. (1920), *The Economics of Welfare*, 4th edn, MacMillan, London.
- Stancil, P. (2004), 'Still Crazy After all These Years: Understanding the Robinson-Patman Act Today', *American Bar Association Section of Business Law* **14**(1), 39–44.
- Varian, H. R. (1989), Price Discrimination, in R. Schmalensee & R. D. Willig, eds, 'Handbook of Industrial Organization', Vol. 1, Elsevier Science, Amsterdam, pp. 597–654.